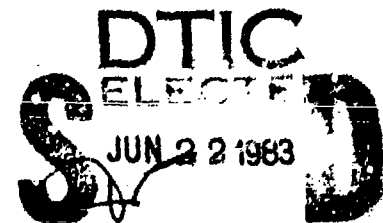


ADA 1 29678

**SUPERVISORY CONTROL OF
REMOTE MANIPULATORS,
VEHICLES AND DYNAMIC
PROCESSES: EXPERIMENTS
IN COMMAND AND DISPLAY
AIDING**

Thomas B. Sheridan

March, 1983



CONTRACT N00014-77-c-0255
WORK UNIT NR 196-182
ENGINEERING PSYCHOLOGY PROGRAMS
OFFICE OF NAVAL RESEARCH
ARLINGTON, VIRGINIA 22217

Approved for public release. Distribution
unlimited. Reproduction in whole or in
part is permitted for any purpose of the
United States Government.

FILE COPY

83 06 22 015

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
	AD-A129678	
4. TITLE (and Subtitle) Supervisory Control of Remote Manipulators, Vehicles and Dynamic Processes: Experiments in Command and Display Aiding		5. TYPE OF REPORT & PERIOD COVERED Final Report
7. AUTHOR(s) Thomas B. Sheridan		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Massachusetts Institute of Technology Cambridge, MA 02139		8. CONTRACT OR GRANT NUMBER(s) N00014-77-C-0256
11. CONTROLLING OFFICE NAME AND ADDRESS Engineering Psychology Group Office of Naval Research, (Code 442EP) Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 196-152
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE 28 February 1983
		13. NUMBER OF PAGES 99
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES none		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) <div style="display: flex; justify-content: space-between;"> <div> supervisory control man-machine systems human performance man-computer interaction </div> <div> teleoperators remote control remote undersea vehicles unmanned vehicles </div> <div> robotics manual control human factors computer control </div> </div>		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report is about supervisory control, an increasingly prevalent form of man-machine system wherein a human operator controls a process as the supervisor of a computer. The computer, in turn, may perform limited automatic control or it may process and display information from sensors. The particular context of interest here is supervisory control of manipulators and vehicles for remote inspection and work in the deep ocean.		

CONFIDENTIAL

CONFIDENTIAL

After giving a more detailed definition of supervisory control and providing examples, the report reviews a number of experimental studies conducted recently at the MIT Man-Machine System Laboratory. These are divided into two groups. The first group of studies is concerned with computer mediation in command and control of manipulation. Two developed systems are described by which an operator may "teach" a manipulator to perform simple manipulation tasks. Other experiments relate to the special problems of communicating geometric information, to compensating for motion disturbances, and to the operator's sometime dilemma between allocating an automatic control system vs. doing the task himself.

The next group of experimental studies examines computer mediation in processing sensed information and displaying it to the human supervisor. The first experiment deals with effects on manipulator control of tradeoffs between frame-rate, resolution and grayscale under severe bandwidth constraints. Three subsequent experiments treat the use of computer-generated models to aid planning and real-time control under conditions of limited feedback or time delay. Two final experiments are concerned with aiding the operator in detecting and locating failures.

A brief conclusion reviews how these experiments fit together and speculates on problems and prospects for supervisory control in manipulation, vehicle and process control, and other areas.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	



I. DEFINITION AND EXAMPLES*

A. What Is Supervisory Control

1. Definition

Supervisory control of a process means a human operator communicates with a computer to gain information and issue commands, while the computer, through artificial sensors and actuators, implements these commands to control the process. Thus a restrictive use of the term supervisory control means that one or several human operators are setting initial conditions for, intermittently adjusting, and receiving information from a computer which itself closes a control loop through external sensors, effectors and the task environment. We may call this definition A.

However, when the computer makes a sufficiently complex transformation of environmental data to produce integrated (chunked) displays, and/or retransforms operator commands to implement sufficiently detailed control actions - even though there is not intermediate loop closure - one may call this supervisory control also (definition B). The essential difference of (B) from (A) is that in (B) the computer cannot act on new information without new authorization by the supervisor (i.e., the computer implements discrete sets of instruction open loop), though the two cases may look similar to the supervisor who still sees and acts through his computer (analogous to his staff). That is, the supervisor may not know whether his subordinates act open-loop or closed-loop.

Any process can be brought under human supervisory control and thus be subsumed under this definition, including vehicles (aircraft, spacecraft, ships, submarines, ground vehicles of all kinds), continuous product processes (oil, chemical, fossil and nuclear power plants), discrete product processes (manufacturing, construction, farming), robotic/teleoperator devices where not included above, and information processing of all kinds (air traffic, military command and control, office automation, etc.).

Figure 1 characterizes supervisory control in relation to manual control and automatic control. Common to the five man-machine system diagrams are displays and controls interfaced with the human operator, and sensors and actuators interacting with a process or "task". The first two systems on the left represent manual control. (1) is without computer aiding while in (2) significant computer transforming or aiding is done in either or both sensing and acting (controlling) loops. Note that in both (1) and (2) all control decisions depend upon the human operator. When either the minor (3) or major (4) fraction of control is accomplished by control loops closed directly through the computer we call this supervisory control. If, once the control system is set up, essentially all the control is automatic (5), that is, if the human operator can observe but cannot influence the process (other than pulling the plug), it is no longer supervisory control.

The five diagrams are ordered with respect to degree of automation. The progression is not meant to imply either degree of sophistication or degree of desirability.

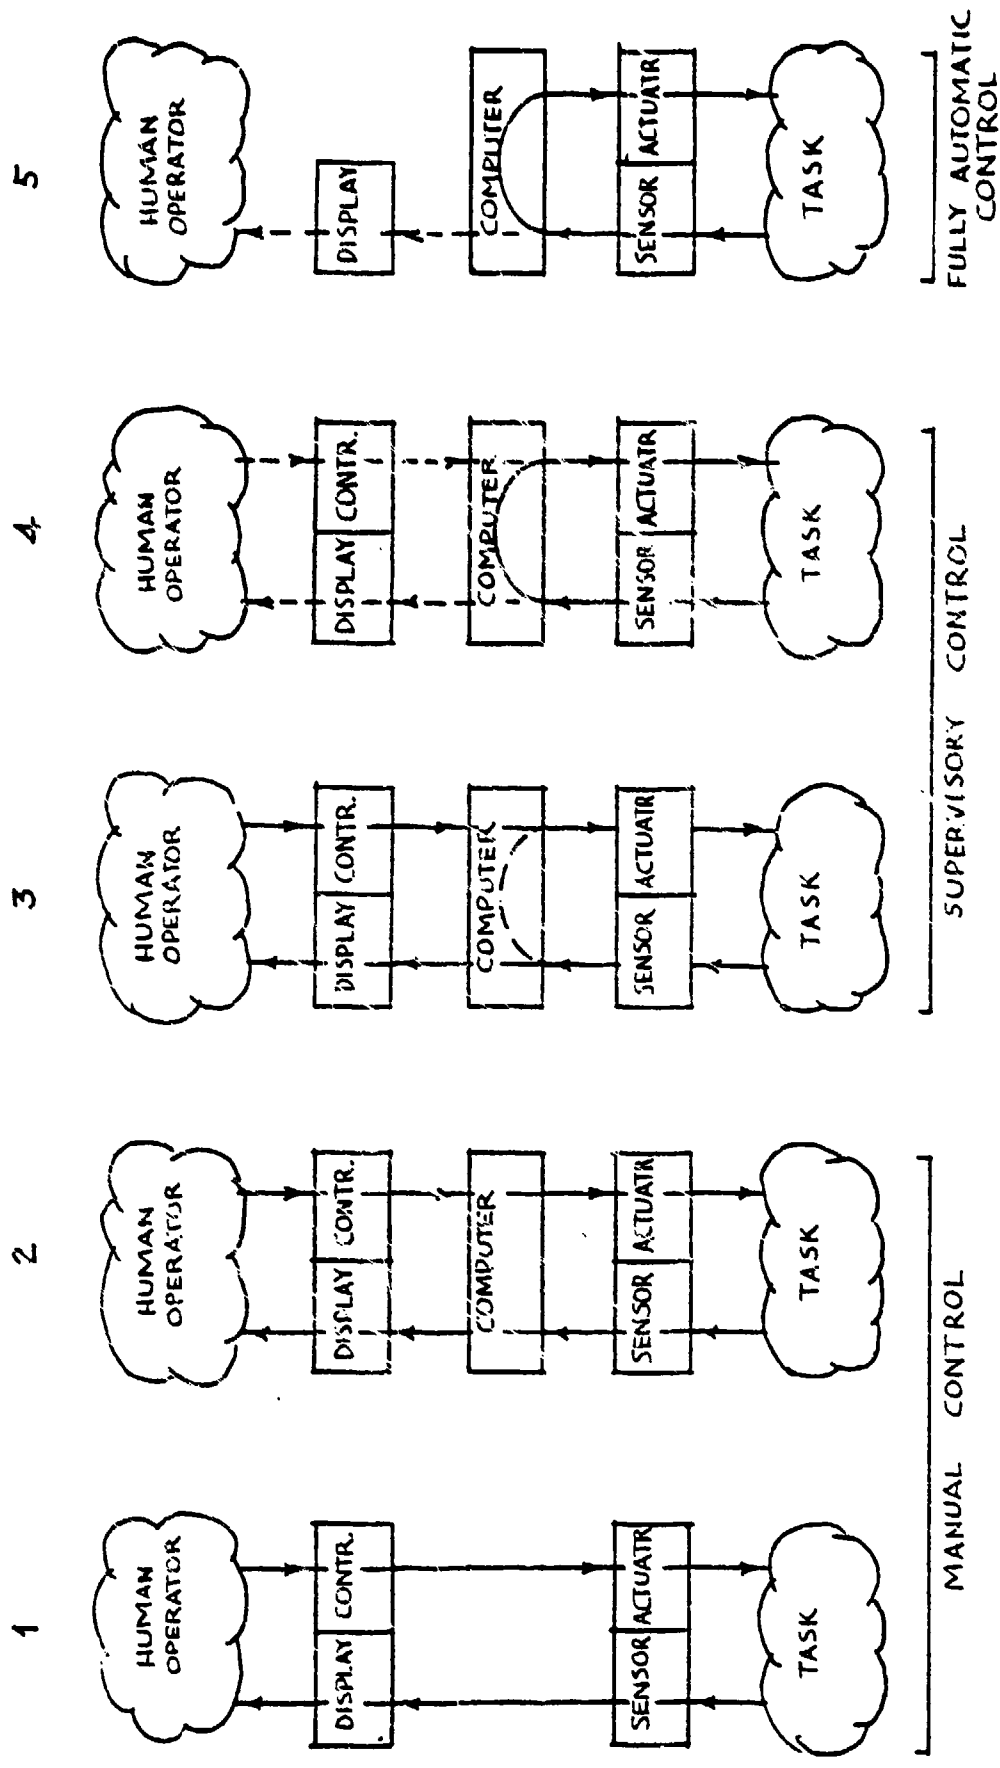


Figure 1. Five man-machine systems ordered as to degree of automation

2. A Model

Figure 2 shows a more general model of a supervisory control system than Figure 1. The human component is still left as a single entity. There are two subsystems, the human-interactive subsystem (HIS) and the task-interactive subsystem (TIS). The HIS generates requests for information from the TIS and issues high level commands to the TIS (subgoal statements, instructions on how to reach each subgoal or what to do otherwise, and changes in parameters). The TIS, insofar as it has subgoals to reach, instructions on how to try or what to do if it is impeded, functions as an automaton. It uses its own artificial sensors and actuators to close the loop through the environment and do what is commanded.

Note that the HIS and TIS form mirror images of one another. In each case the computer closes a loop through mechanical displacement (hand control, actuator) and electro-optical or sonic (display, sensor) transducers to interact with an external dynamic process (human operator, task). The external process is quite variable in time and space and somewhat unpredictable.

The numbered arrows identify individual cause-effect functions, with explanations of the loops at the right. It is seen that there are three types of inputs into the human operator: (1) those which come by loop 1 directly from the task (direct seeing, hearing or touching), (2) those which come by loops 2 and 8 through the artificial display and are generated by the computer and (3) those which come by loops 10 and 9 from the display or manual controls without going through the computer (i.e. information about the display itself such as brightness or format, present position of manual controls, which is not information which the computer has to provide). Similarly, there are three types of human outputs: (1) those which go by loop 6 directly to the task (the human operator by-passes the manual controls and computer and directly manipulates the task, makes repairs etc.) (2) those which communicate instructions via loops 7 and 8 to the computer, and (3) those which modify the display or manual control parameters via loops 10 and 9 without affecting the computer (i.e., change the location, forces, labels or other properties of the display or manual control devices).

Correspondingly there are three types of force and displacement input into the task: (1) direct manipulations by the operator via loop 6; (2) manipulations controlled by the computer via loops 3 and 7; and (3) those forces which occur by interaction, over loops 4 and 5, with the sensors and actuators and are not mediated or usually what was intended by the computer or operator. Finally there are three types of outputs from the task: (1) information fed back directly to the operator over loop 1; (2) information fed to the TIS computer via loops 2 and 3; and (3) information (in the form of forces and displacements) which modifies the sensors or actuators via loops 4 and 5 without being explicitly sensed by the computer.

When the task is near to the operator, the HIS and TIS computers can be one and the same. When the TIS is remote usually HIS and TIS computers are separated to avoid problems caused by bandwidth or reliability constraints in telecommunication, loops 2 and 7. This problem will be discussed in detail in section IV-C.

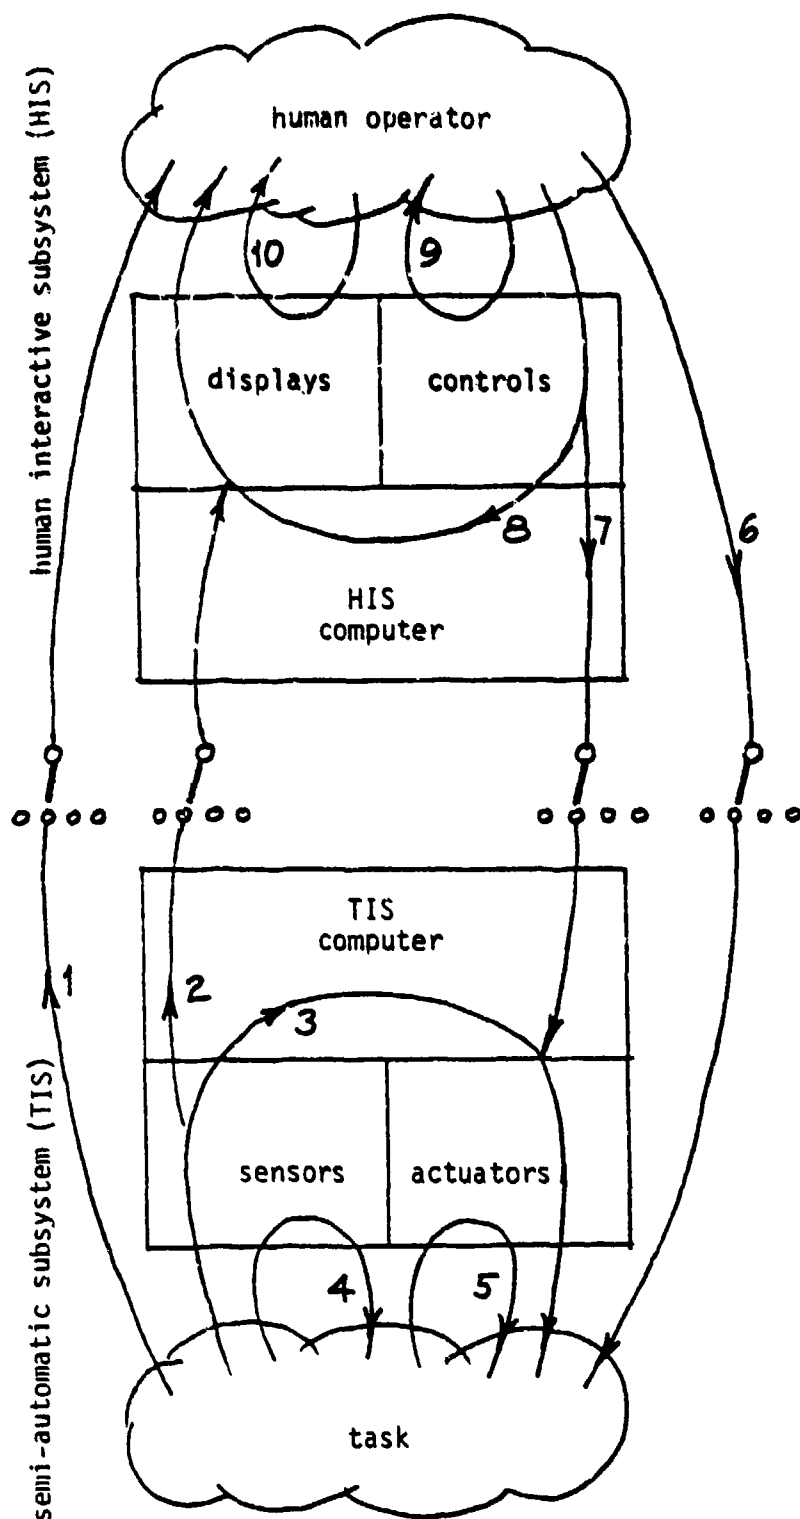


Figure 2. General model of supervisory control system

Multiplexing switches are shown in loops 1,2,7, and 6 to suggest that one HIS may be time-shared among many TIS, i.e., many tasks, each with its own local automatic control or robotic implementer. In fact, more and more this is coming to be the case in supervisory control. In some process plants there are over 1000 TIS, some being primitive feedback controllers, some being simply programmed but highly reliable microcomputers. The sheer number of TIS causes a multiplexing or switching overhead cost.

The above is a descriptive model of supervisory control; that is, it is intended to fit what is observed to be the structural and functional nature of a wide variety of situations we discussed earlier. The variables on the lines of Figure 1 are all measurable; there are no intervening variables, no suppositions about what is going on that we cannot observe readily.

Being a descriptive model this is by definition not a normative model. We have not imposed any notions of how the system should work, or of what optimal behavior consists, or how close actual behavior compares to optimal.

It is important to note, also, that we do not intend to develop a model of the human operator independent of the rest of the system. McRuer and Krendel [1965] abandoned trying to model the human operator in a simple control loop as an invariant entity per se and turned instead to finding invariance in the series combination of human controller plus controlled process. It seems that it is best to find invariance in supervisory control plus tool box of computers, sensors and effectors plus task. One may note various functions that either human or computer can do, but some are best done by human and some are best done by computer. Which does what function will evolve over many years in the future and will always depend on circumstance. For now the intent is to provide a qualitative description of the combination.

The essence of supervision, as noted in conjunction with the dictionary definition earlier, is that it is not a single activity, as we are accustomed to characterize various sensory-motor or cognitive or decision-making skills, or communication or controlling behavior. Supervision implies that the primary or direct activity, whatever it is, is normally being done by some entity (man or machine) other than the supervisor. There may be a single primary task, or many such tasks. The supervisor from outside, performs those many functions necessary to insure that the single entity does (or multiple entities do) what he, the supervisor, intends. Thus there may be multiplicity of function for two reasons:

1. For each primary task there are many different things to do to ensure that the primary entity (what was called the TIS in the model description) does what the supervisor intends that it do.
2. When there are multiple primary tasks, while the basic functions may be similar from one primary task to another, the data are different, and the initial conditions are different in performing the same function on each.

Our supervisory control model shows the supervisory computer to

multiplex among or alternately connect to, different TIS's or primary tasks. It also shows multiple connections to and from the human operator. It does not make clear that in switching from one TIS to another the initial conditions ("getting one's bearings") are different with each switch. Nor does it make clear that the human supervisor is continually switching functions even while dealing with a single TIS.

But this seems to be the essence of the supervisor: switching functions within one task and switching tasks. The remainder of this section elaborates this point.

Earlier, in conjunction with the dictionary definition, the ideas of "planning," "programming" and "observing" emerged as different components of "supervising".

Missing explicitly from the earlier dictionary definition but implied nevertheless are two additional functions. The first is taking over from the "other" entity, the TIS in our case, seizing direct control when indirect control by supervision fails. The second is to learn from experience.

Summarizing and elaborating to suit our present context, the supervisor, with respect to each task (and each TIS), must perform five distinct functions listed in Table 1. While the explanation of these five functions in Table 1 is not a consensus and some of these steps are manifest to a greater or lesser degree in any actual supervisory control situation, the point is that the necessary sequencing through these differing functions makes the human supervisory controller essentially different from the human in-the-loop, one-continuous function controller and decision-maker.

As implied above the allocation of attention by the human supervisor is both between functions for a given tasks and between tasks. In skilled or overlearned activities a person can engage in many at once (provided the required sensors and effectors are not overtaxed with respect to simple mechanical or signal processing considerations). Thus one can drive a car, talk, scratch his nose and look for a landmark at the same time. But one cannot do multiple simultaneous tasks each of which requires "new thinking" unless the time requirements are such that one can shift attention back and forth. In view of these facts we initially may characterize the attention allocation of the human supervisor as, first, selecting among alternative tasks to be done, and second, selecting his proper function with respect to that task.

B. Examples Of Supervisory Control In Remote Manipulation

Intrinsic to supervisory control is the idea of teleoperation - man performs a sensing and/or manipulation task remotely by use of artificial sensors and actuators. This can be spatial remoteness, as with a remotely controlled vehicle or manipulator undersea or in space. It can be temporal remoteness, due to a time delay between when an operator issues commands and when he receives feedback. Or it can be functional remoteness, meaning that what the operator sees and does and what the system senses and does bear little superficial resemblance. Teleoperation can be either the motivation for or the result of supervisory control, as will be made evident.

Table 1. Functions of the supervisor

<p>1. Plan</p> <ul style="list-style-type: none"> a) be aware of what tasks are to be done, what resources are available, what resources (TIS) are committed to what tasks, and what resources are uncommitted b) decide on overall goal or goals, including objective function or tradeoffs among goals, and including criteria for handling uncertainties c) decide on strategy or general procedure, including logic of authority (human, HIS computer, TIS computer) in various situations d) consider known initial conditions and various combinations of probable inputs and possible actions and their consequences in view of system constraints and capabilities e) determine best action sequence to do what is intended under various situations f) decide what is to be considered abnormal behavior including automatic recovery from trouble, and what defaults or contingency actions are appropriate.
<p>2. Teach (a, b, c and d could also be considered part of planning)</p> <ul style="list-style-type: none"> a) estimate what the computers (HIS and TIS) know of the situation b) decide how to instruct the HIS to instruct the TIS to execute intended and abnormal actions c) decide how many of intended and abnormal actions TIS should undertake in one frame, i.e. before further instruction d) try out part or all of that instruction in (operator's) own mental and/or HIS computer model without commitment to transmit to TIS e) impart instruction (program) to HIS computer with commitment to transmit to TIS f) give command to HIS to start action
<p>3. Monitor</p> <ul style="list-style-type: none"> a) decide on what TIS behavior to observe b) specify to HIS computer the desired display format c) observe display, looking for signals of abnormal behavior and performing on-line computer-aided analysis of trends or prediction or cross-correlation as required d) observe task directly when and if necessary e) make minor adjustments of system parameters when necessary, as the automatic control continues. f) diagnose apparent abnormalities or failure, if they occur, using computer aids
<p>4. Intervene</p> <ul style="list-style-type: none"> a) decide when continuation of automatic control would cease to be satisfactory and minor parameter adjustments would not suffice either b) go physically to TIS or bypass all or portions of HIS and TIS computers to effect alternative control actions or stoppage or recovery c) implement maintenance or repair or modifications of TIS or task d) recycle to (1), (2) or (3) as appropriate
<p>5. Learn</p> <ul style="list-style-type: none"> a) decide means for collecting salient data and drawing inferences from it over repeated system runs b) implement these means c) allow for serendipitous learning d) periodically take stock of learning, modify system hardware and software, and anticipate future planning of operations e) develop understanding of and trust in the system

In a sense, manipulators combine the functions of process control and vehicle control. The manipulator base may be carried on a spacecraft, a ground vehicle, or a submarine, or its base may be carried on a spacecraft, a ground vehicle, or a submarine, or its base may be fixed. The hand (gripper, end effector) is moved relative to the base in up to three degrees of translation and three degrees of rotation. It may have one degree of freedom for gripping, but some hands have differentially movable fingers or otherwise have more degrees of freedom to perform special cutting, drilling, finishing, cleaning, welding, paint spraying, sensing, or other functions.

Manipulators are being used in many different applications, including lunar moving vehicles, undersea operations, and hazardous operations in industry. The type of supervisory control and its justification differs according to the application.

The fact of a three-second time delay in the earth-lunar control loop resulting from round-trip radio transmission from earth leads to instabilities, unless an operator waits three seconds after each of a series of incremental movements. This makes direct manual control time-consuming and impractical. Sheridan and Ferrell [1967] proposed having a computer on the moon receive commands to complete segments of a movement task locally using local sensors and local computer program control. They proposed calling this mode supervisory control. Delays in sending the task segments from earth to moon would be unimportant, so long as rapid local control could introduce actions to deal with obstacles or perform other self-protection functions.

The importance of supervisory control to the undersea vehicle manipulator is also compelling. There are things the operator cannot sense or can sense only with great difficulty and time delay (e.g., the mud may easily be stirred up, producing turbid opaque water that prevents the video camera from seeing), so that local sensing and quick response may be more reliable. For monotonous tasks (e.g., inspecting pipelines, structures, or ship hulls or surveying the ocean bottom to find some object) the operator cannot remain alert for long; if adequate artificial sensors could be provided for the key variables, supervisory control should be much more reliable. The human operator may have other things to do, so that supervisory control would facilitate periodic checks to update the computer program or help the remote device get out of trouble. A final reason for supervisory control, and often the most acceptable, is that, if communications, power, or other systems fail, there are fail-safe control modes into which the remote system reverts to get the vehicle back to the surface or otherwise render it recoverable.

Many of these same reasons for supervisory control apply to other uses of manipulators. Probably the greatest current interest in manipulators is for manufacturing (so-called industrial robots), including machining, welding, paint spraying, heat treatment, surface cleaning, bin picking, parts feeding for punch presses, handling between transfer lines, assembly, inspection, loading and unloading finished units, and warehousing. Today repetitive tasks such as welding and paint spraying can be programmed by the supervisor, then implemented with the control loops that report position and velocity. If the parts conveyor is sufficiently reliable, welding or painting nonexistent objects seldom occurs, so that more sophisticated feedback, involving touch or vision, is usually not required. Manufacturing

assembly, however, has proven to be a far more difficult task.

In contrast to assembly line operations, in which, even if there is a mix of products, every task is prespecified, in many new applications of manipulators with supervisory control, each new task is unpredictable to a considerable extent. Some examples are mining, earth moving, building construction, building and street cleaning and maintenance, trash collection, logging, and crop harvesting, in which large forces and power must be applied to external objects. The human operator is necessary to program or otherwise guide the manipulator in some degrees of freedom, to accommodate each new situation; in other respects certain characteristic motions are preprogrammed and need only to be initiated at the correct time. In some medical applications, such as microsurgery, the goal is to minify rather than enlarge motions and forces, to extend the surgeon's hand tools through tiny body cavities to cut, to obtain tissue samples, to remove unhealthy tissue, or to stitch. Again, the surgeon controls some degrees of freedom (e.g., of an optical probe or a cauterizing snare), while automation controls other variables (e.g., air or water pressure).

II. ILLUSTRATIVE EXPERIMENTS ON COMMAND AIDS TO THE SUPERVISOR

The experiments described below are concerned with the "effector" loops of Figure 2 (loops 6 and 7) and the automation (loop 3) which augments operator control. Sections II-A and II-B describe the development of particular supervisory command/control systems for remote manipulation. Section II-C examines the problems of "pointing", that is, telling the computer about the specific geometrical positions and orientations in the world which it must know in order to do its job; that such nominally motor activity will be shown to be highly influenced by perceptual factors points up the artificiality of trying to separate effector from sensory mechanisms. Section II-D exemplifies a simple form of automatic error nulling within a manual loop, a primitive form of supervisory control. Section II-E discusses some relatively abstract laboratory research dealing with a new dilemma facing the supervisor - when should he allocate a machine to do a task and when should he do it himself.

A. Supervisory Command Of Remote Manipulation: "SUPERMAN"

* Portions of this section are from Brooks and Sheridan [1980]

1. Introduction

To investigate the relative merits of supervisory control applied to teleoperators, specifically telemanipulators in this case, a task-referenced sensor-aided supervisory system, called SUPERMAN, was built and experiments were performed. These experiments compare various conventional control modes with supervisory control, and demonstrate that supervisory manipulation does improve performance in the majority of cases.

2. Method and Apparatus

The major elements of the SUPERMAN system are a modified Argonne E2 master-slave manipulator with six degrees-of-freedom, a dedicated control interface (DASI), and an Interdata 70 computer. Designed for efficient man-machine interaction with both analog and symbolic control inputs, the system can be commanded by a variety of conventional control modes as well as supervisory. In addition, time delay and/or noise can be added for experimental purposes.

Using both analog and symbolic commands, a manipulation can be taught and/or demonstrated to the computer. Trained manipulations can be transformed from one coordinate system to another so that once the generic characteristics of a task have been learned, the machine can perform similar tasks in different locations without further training. When the human operator requires a particular trained manipulation he simply "initializes" the new coordinate system relative to the old by moving the teleoperator hand to the starting point of the task (e.g., grasping a nut or valve handle) and signals for execution. Certain objects in the task environment can, of course, maintain their original coordinates.

Since the E-2 manipulator can sense the forces generated during the task, supervisory programs can call for repeated movements which, upon certain touch conditions becoming true, branch into other movements. For example, repeated hand movements can grasp a nut, unscrew it by one

revolution, pull repeated hand movements and grasp a nut, unscrew it by one revolution, pull back to test whether it is off and, if it is, place it in a bucket or, if it is not, repeat the operation. Similar supervisory programs have been applied to attaching a nut to a bolt, opening and closing a valve, scooping dirt and so on. Further information on the SUPERMAN system can be found in Brooks [1980].

The manipulator laboratory was arranged as shown in Figure 3a during the experiments. To simulate remote conditions the operator viewed the task environment through either a mono or 2-view television system. The video system consisted of two black and white high-resolution 9 in. monitors, a fixed camera with wide angle lens, and a zoom camera with pan and tilt.

Figure 3b shows the manipulator environment and the experimental tasks designed for this study. The tool rack and sample buckets remained in the locations shown throughout the experiments since these pieces of equipment are usually rigidly attached to the teleoperator vehicle in real applications. Also shown in the figure are the movable task hub and task board on which representative tasks such as valves, bolts, etc. were mounted. The location of the task hub and board were changed throughout the study to simulate the random task/vehicle relationships which are typical of the arbitrary environments found in marine and space applications.

3. Experimental Design

Six manipulation tasks were identified for Experimental Investigation: (1) tool retrieval; (2) tool return; (3) taking a nut off; (4) grasping an object and placing it in a container; (5) opening/closing a valve; and (6) digging. In addition, four manual control modes were delineated as important experimental parameters: (1) switch fixed rate; (2) joystick variable rate; (3) master-slave position control; and (4) master-slave position control with force reflection. With regard to the video arrangement, both mono and 2-view conditions were tested for comparison. Due to time constraints only three subjects were used for four of the tasks (tool retrieval, tool return, nut-off and sampler), and only one subject was used for the remaining two (open/close valve and digger). Each experiment was performed 5 times by each subject to obtain a statistical mean and standard deviation. Both manual and supervisory control were used.

These conditions result in a total of 1120 experimental runs. Since this would require an inordinate amount of time, the experimental load was reduced to 680 runs by noting that some of the tasks, or portions of the tasks, had constant computer execution times.

4. Subjects and Training. Three classes of subjects were used for these experiments, one experienced, four well trained, and two untrained subjects. The experienced subject had over 200 hours of training of this particular system. The well trained subjects had an average of 20 hours training given in 15 minute intervals for each of the control modes.

5. Procedure. The experiments were scored on the basis of recorded time and errors. The subjects were not given specific instructions to minimize either quality, but only to weigh them equally.

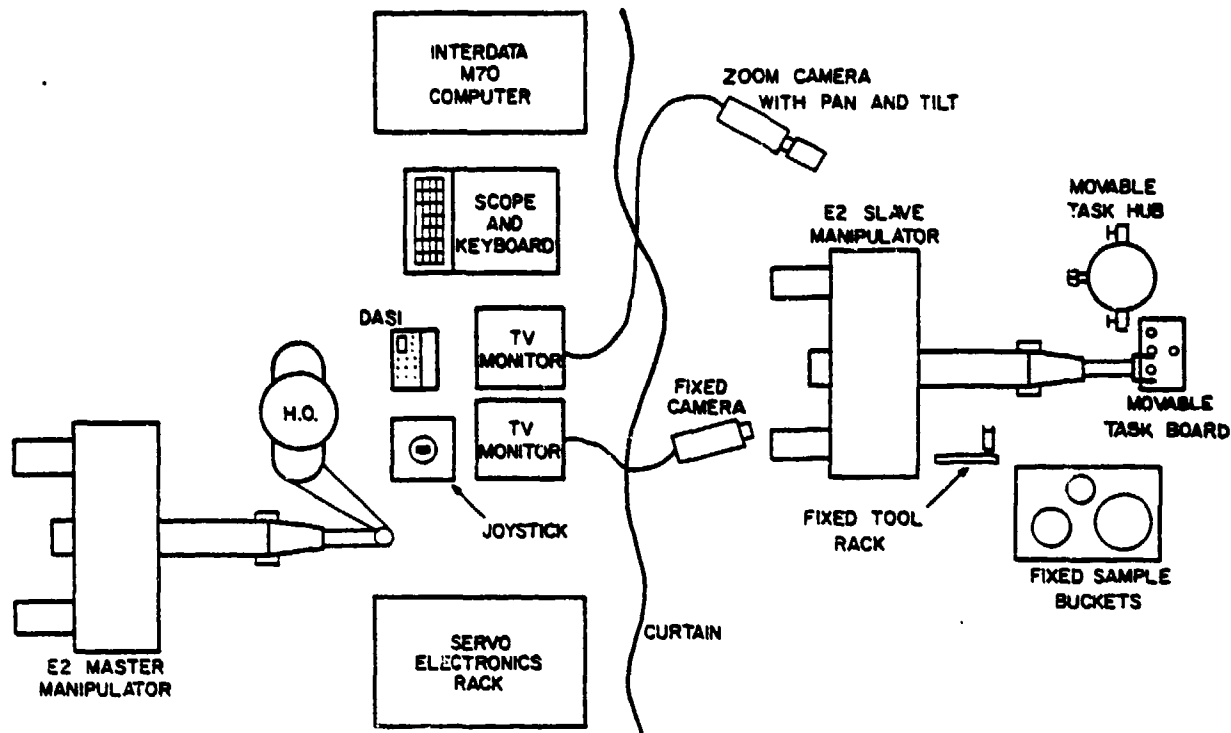


Figure 3a : Schematic of Experimental Layout

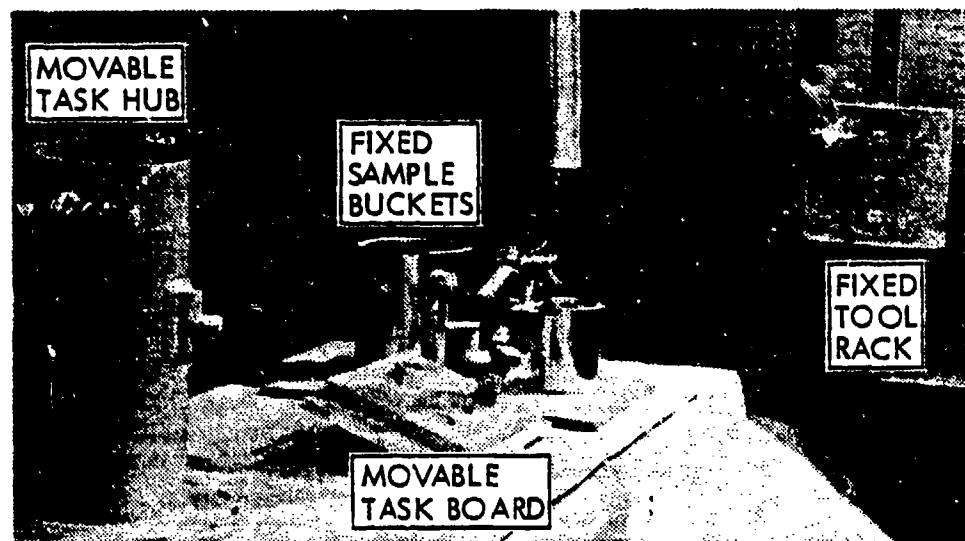


Figure 3b : Task Hub, Task Board, Sample Buckets, and Tool Rack

Tool-Retrieval Task - The first task required the subject to start with the end effector positioned near the task hub. On the experimenter's signal, the subject moved the end effector to the tool rack, obtained the tool, being sure it was properly seated in the hand, and returned with the tool to the starting position. The subjects were told that the success or failure of the task was measured by whether a solid connection between the tool handle and end effector was achieved. Execution of this task under supervisory control simply involved a button push.

Tool-Return Task - For the second task the subject started from a position next to the task hub with the tool in hand, and on the experimenter's signal, moved to the rack, replaced the tool insuring that it was properly seated, and returned to the initial position. The operators were told that the success or failure of the task was determined by whether or not the tool was properly replaced on the rack. To properly seat the tool on the rack required that both of the 1/8 inch rack pins were engaged in the handle and that the tool was completely pushed onto the pins. This task was executed under supervisory control through a simple button push.

Nut-Removal Task - This experiment began with the end effector positioned over the valve on the task hub. On the experimenter's signal, the subject moved the end effector from the valve to the nut, oriented the hand, and removed the nut. The general procedure used by the subjects and computer was to turn 180°, pull back to test if the nut was off, and then either reverse 180° and continue, or remove the nut. Prior to the task, the operators were told that the task would be considered successfully completed if the nut could be removed without losing it. Under supervisory control the operator initialized the task by moving from the starting position to the nut, orienting the hand with the rotational axis of the nut, and signaling the computer to remove it.

Sampling Task - The fourth task required the subject to pick-up thirteen randomly placed samples and put them in one of two buckets according to their size. The subjects were told that their success or failure to complete the task would be measured by how many samples were successfully placed in the proper buckets. Under supervisory control the operator initialized the task by placing the end effector over the sample and signaling the computer to place it in the appropriate bucket. The computer returned control to the subject at the location where the sample was grasped. The operator then moved to another sample, initialized, and continued until all 13 samples were in the buckets.

Open/Close Valve Test - This experiment required the subject to position the end effector over the nut on the task hub, and then, on the experimenter's signal, the subject moved to the valve, oriented the hand, and opened or closed the valve as required (opening and closing tasks were switched after each experiment). The subject was required to continue until the valve operation was complete. To initialize this task under supervisory control the operator oriented the end effector on the rotational axis of the valve and signaled the computer either to open or close it as required. The computer checked the rotational torques to determine if the task had been completed.

Digging Task - The final task required the subject to remove a

specified amount of soil from a box by filling a bucket with a shovel. This task is composed of a number of subtasks: (1) the shovel is positioned to remove the soil, (2) the shovel is pushed into the soil and lifted out, and (3) the soil is transported to the bucket and dropped in. The subject was required to continue until the bucket was filled. Under supervisory control the positioning of the shovel was performed manually (i.e., the operator decided when and where to dig) while the scooping and dropping actions were executed by the computer.

6. Results

It has been shown by a number of investigators that the time required to perform a task can be attributed to a number of distinctly different motions. For example, one classification divides the task time for control with a time delay into segments related to get, transport, and position motions. For a peg-in-the-hole task Hill [1976] has shown that there are two independent motions which determine the total task time under manual control - gross travel and precision. This report will use a similar scheme to describe the task completion time for a supervisory system:

$$t_{TT} = t_I + t_P$$

where

$$\begin{aligned} t_{TT} &= \text{Task Time} \\ t_I &= \text{Time required by the human operator to initialize the task. This time is primarily a function of the initial hand/task locations and the manual control mode used to locate the task.} \\ t_P &= \text{Time required by the computer to perform the task. This time is primarily a function of the task complexity.} \end{aligned}$$

The determination of these times is rather simple due to the discontinuity in control which occurs during the trade from manual initialization to computer execution (this "discontinuity" is a desired result since trading of control should be "apparent").

Figures 4-7 are plots of typical data. The data recorded during the supervisory experiments have been divided into initialization and performance times to indicate the time spent by each action. Each of the time bars is the result of data averaged over two trained subjects, except for Figure 7 which is averaged over three trained subjects. The lines to the left of the manual control bars give the range over which the trained subjects performed the task. For comparison, the average time for an inexperienced subject to perform the first three tasks is also given (denoted by triangles). The mean times of the untrained subjects were always above the maximum value of the trained subjects. The lines to the left of the manual control bars give the range over which the trained subjects performed the task. For comparison, the average time for an inexperienced subject to perform the first three tasks is also given (denoted by triangles). The mean times of the untrained subjects were always above the maximum value of the trained subjects for the same task and control mode. The lower portion of each Figure (Figure 4b-7b) plots the mean number of errors which occurred under manual and supervisory control.

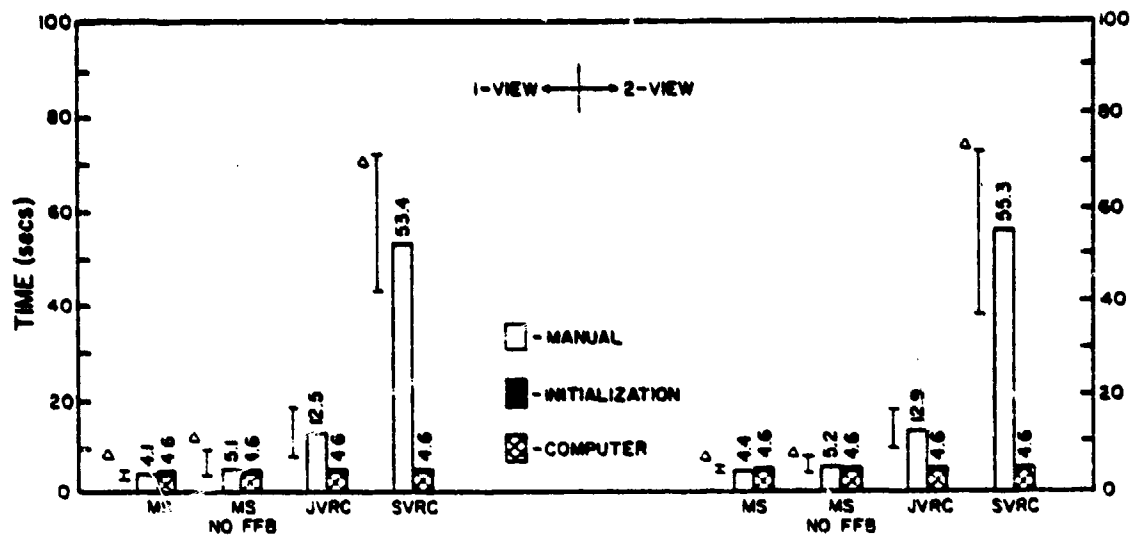


Figure 4a : Average Tool-Retrieval Time. Each bar gives the average time of two subjects. The Δ symbol represents the mean time for an untrained subject. The capped lines show the total range of data for the trained subjects.

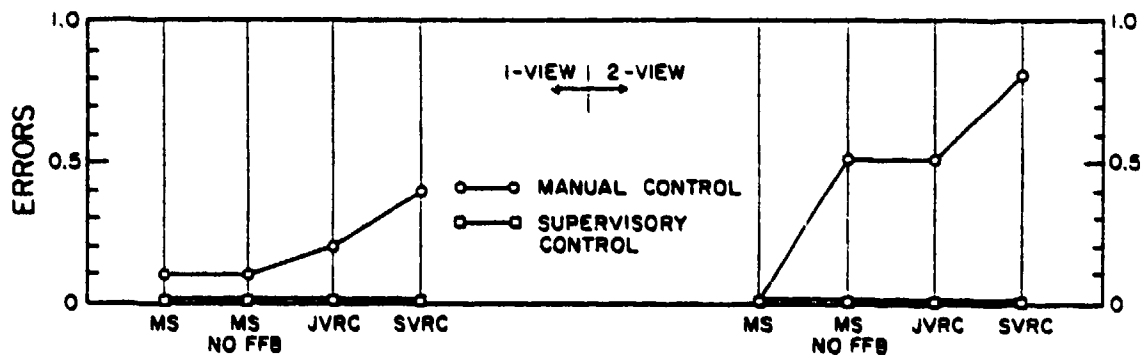


Figure 4b : Expected Number of Tool-Retrieval Errors. Each data point represents the average error rate of two trained subjects. Possible errors included collisions, dropping the tool, and not seating the handle in the end effector properly.

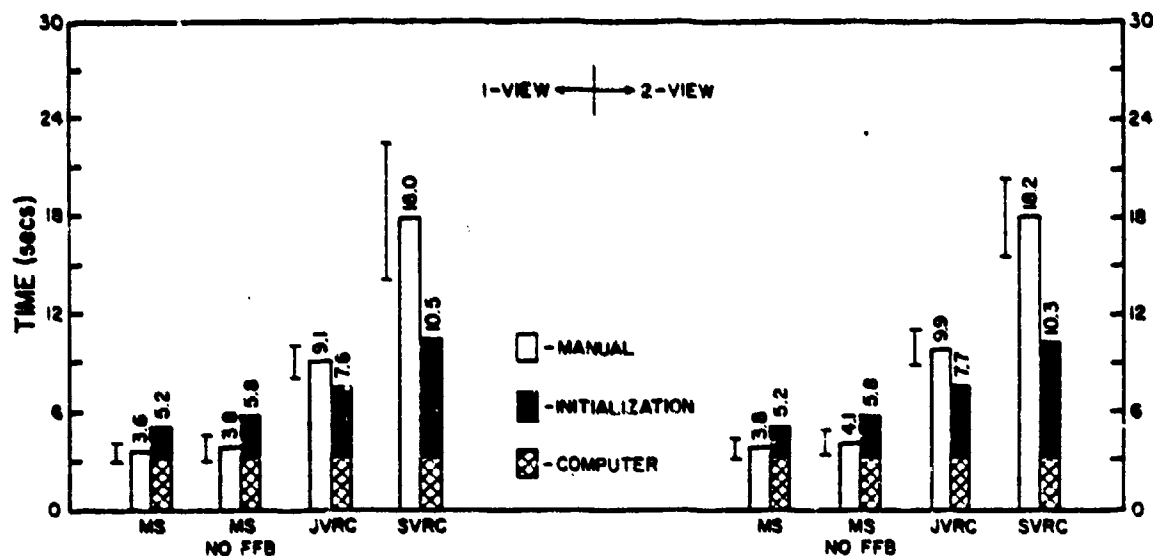


Figure 5a : Average Sampling Time. Each bar represents the mean time of three trained subjects. The capped lines represent the total range of data for the subjects.

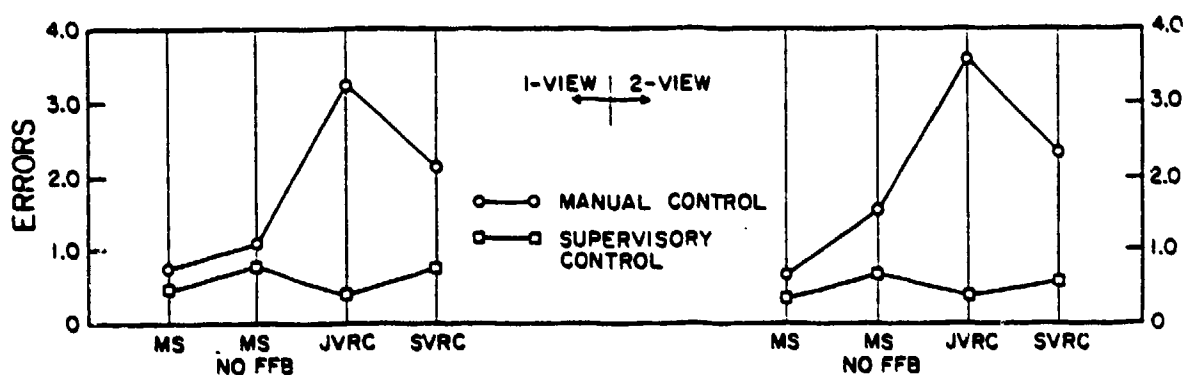


Figure 5b : Expected Number of Sampling Errors. Each data point represents the average error rate of three trained subjects for 13 sampling actions. Possible errors included collisions, missed buckets, lost samples, and (under supervisory control) pressing the wrong button.

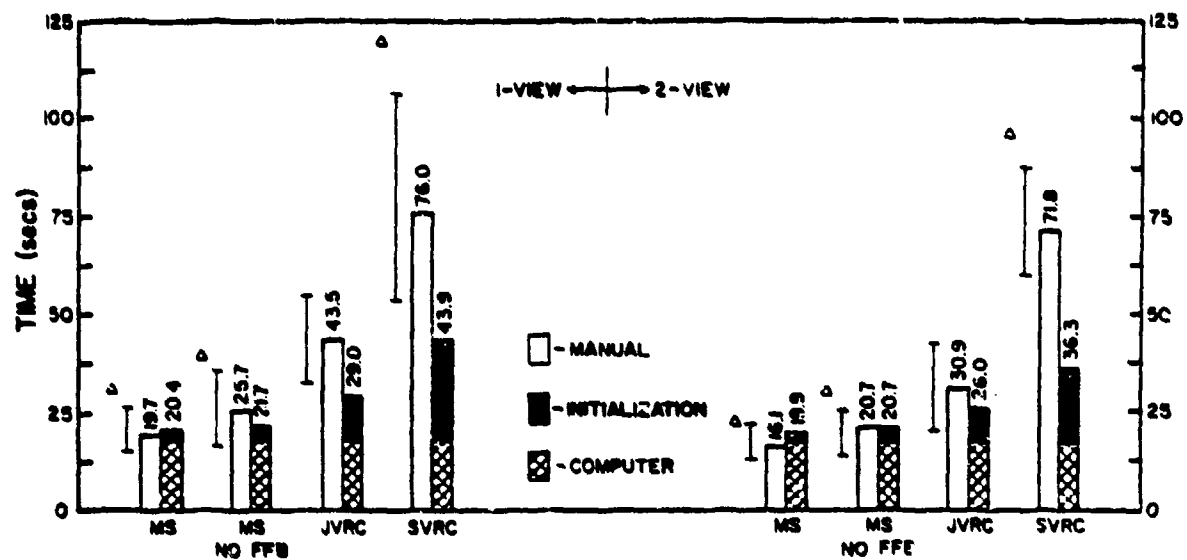


Figure 6a : Average Nut-Removal Time. Each bar represents the average time of two trained subjects and each Δ gives the mean time for an untrained subject. The capped lines represent the total range of data for the trained subject.

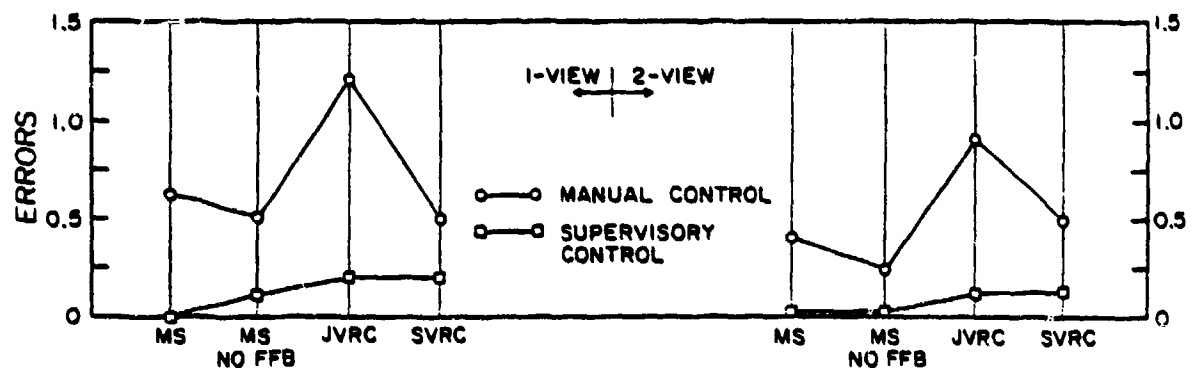


Figure 6b : Expected Number of Nut-Removal Errors. Each data point represents the average error rate of two trained subjects. Possible errors included collisions and dropping the nut.

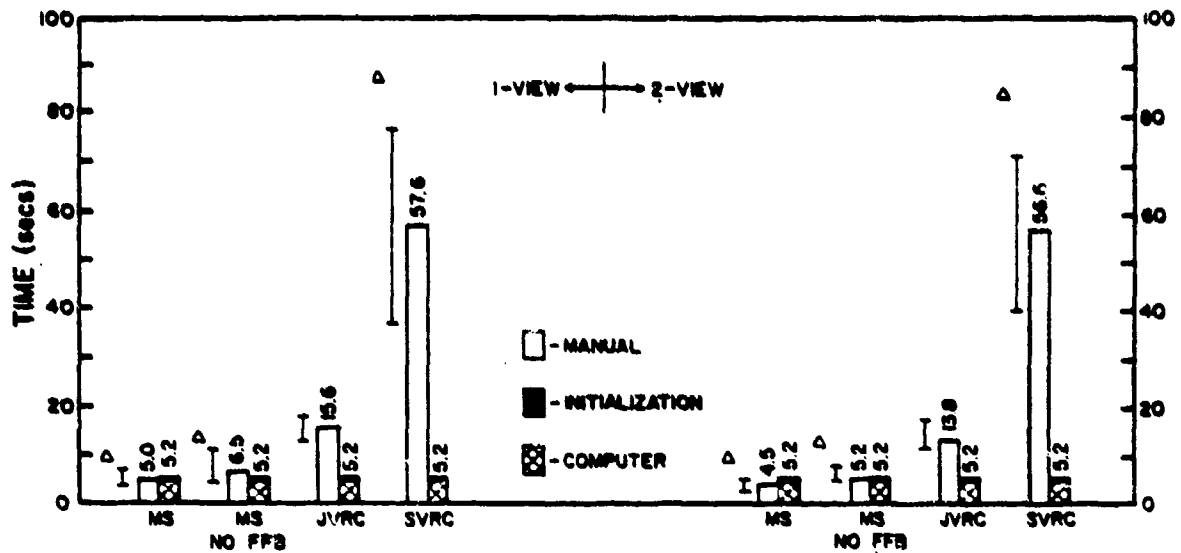


Figure 7a: Average Tool-Return Time. Each bar represents the average time of two trained subjects and each Δ gives the mean time for an untrained subject. The capped lines represent the total range of data for the trained subjects.

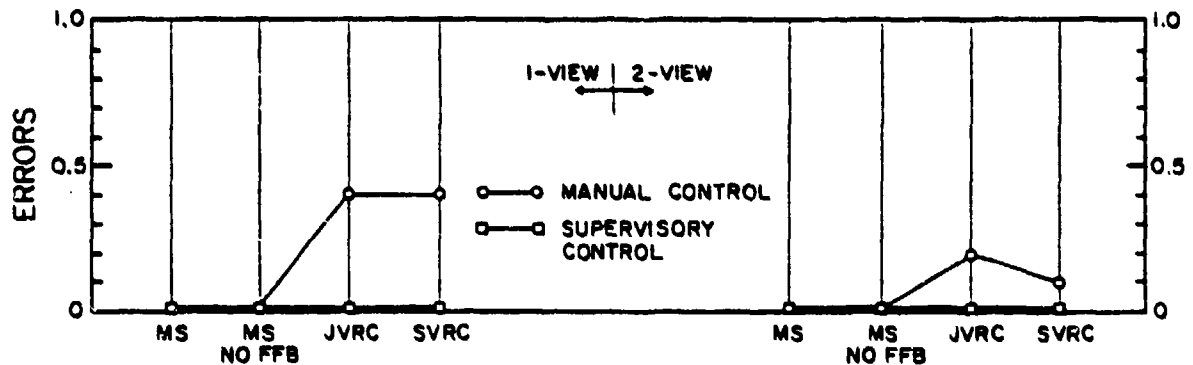


Figure 7b: Expected Number of Tool-Return Errors. Each data point represents the average error rate of two trained subjects. Possible errors included collisions, dropping the tool, and not seating the handle on the rack properly.

7. Manual Control

Predictably, the task completion time increased with control complexity for all tasks. Viewing conditions (mono and 2-view) appeared to affect tasks which required precision movements (e.g., return tool and nut-off), but had little or no effect on the less precise tasks (e.g., sampling). In general, the number of errors increased as the control complexity increased from master-slave to switch rate. However, for some of the tasks a sharp decrease in errors was noticed between joystick and switch rate control (e.g., see Figure 6b and 7b). This effect is attributable to two factors: (1) the increased attention and care each operator exhibited during switch rate control modes (i.e., to move from point A to point B requires considerable thought and effort with switch rate control, but under joystick control the desired movement only requires a push on the stick), and (2) the coincidental matching of the task degrees of freedom and control degrees of freedom (e.g., in the valve or nut-off tasks the axis of rotation corresponded with the hand axis of rotation).

Table 2 gives the ratio of task completion times for each control mode with respect to the "best" control case, master-slave with force feedback. The ratios are given for each subject, task and viewing condition. The untrained subjects are denoted by U1 and U2, the trained subjects are denoted by T1, T2, T3 and T4, and the experienced subject is denoted by E1. The table shows a number of interesting trends: (1) the ratios increase with increasing control complexity, (2) the ratios are approximately constant across subjects (both trained and untrained) within a given task, (3) the ratios are constant across viewing conditions, and (4) the ratios are not constant across tasks (the tasks have been arranged in the table so that the ratio increases as the page is read from top to bottom).

8. Supervisory Control

As would be expected, the time required by the computer to perform its portion of the task remained fixed regardless of the manual control mode from which the human operator issued the execution command. Also, since the only action required of the operator to initiate the tool-retrieval and return tasks was a button push, the absence of initialization times in Figure 4a and 5a was not surprising. The remaining tasks, including those not shown in this report had initialization times associated with the overall task time. As seen in Figure 6 and 7 the initialization times increased with control complexity.

Table 3 gives the ratios of the task completion times under manual control to the times under supervisory control. The ratios are given for each subject, task and viewing condition. The ratios relative to computer control (Table 3) do not show the same trends as those relative to master-slave control (Table 2). It is interesting to note that in contrast to the consistent ratios of Table 2, the computer control ratios of the untrained subjects are significantly higher than the trained subjects: clearly, untrained subjects gain more from supervisory control than trained subjects. Gains from supervisory control for any manual mode are seen to be most significant for tasks which do not require initialization procedures other than a button push (i.e., tool-retrieval and tool-return). The control mode columns clearly indicate the results of the SUPERMAN experiments: (1)

Table 2 : Ratio of Time to Perform Task Under Given Control Mode to Time to Perform Task Under Master-Slave with Force Feedback (CM/MS).

	1-VIEW				2-VIEW				
	MS	MS 70 PER	SVRC	SVRC	MS	MS 70 PER	SVRC	SVRC	
VALVE 1-DOF	1.0	1.2	1.4	1.6	1.0	1.1	1.6	1.9	E1
PUT-OFF 2-DOF	1.0	1.2	2.2	4.3	1.0	1.3	2.2	4.4	T1
	1.0	1.4	2.2	3.3	1.0	1.2	1.9	4.3	E1
	1.0	1.3	*	3.9	1.0	1.4	*	4.3	U2
SAMPLER 3-DOF	1.0	1.1	2.5	5.1	1.0	1.1	2.7	4.4	T2
	1.0	1.0	2.8	5.3	1.0	1.0	2.5	5.1	T3
	1.0	1.0	2.4	4.7	1.0	1.1	2.6	4.9	E1
SCOOPER 7-DOF	1.0	1.1	3.9	10.6	1.0	1.2	3.9	10.9	E1
RETURN-TOOL 6-DOF	1.0	1.1	3.6	10.7	1.0	1.2	3.4	12.1	T4
	1.0	1.4	2.7	12.2	1.0	1.3	2.7	13.2	E1
	1.0	1.4	*	9.0	1.0	1.4	*	9.7	U1
GET-TOOL 6-DOF	1.0	1.1	3.3	13.0	1.0	1.2	3.3	12.1	T4
	1.0	1.4	2.4	13.0	1.0	1.2	2.6	13.2	E1
	1.0	1.4	*	3.3	1.0	1.3	*	11.8	U1

Table 3 : Ratio of Time to Perform Task Under Manual Control to Time to Perform Task Under Supervisory Control (MC/SC).

	1-VIEW				2-VIEW				
	MS	MS NO FFG	SVRC	SVRC	MS	MS NO FFG	SVRC	SVRC	
VALVE 1-DOF	1.1	1.1	1.2	2.0	1.1	0.9	1.1	2.0	E1
WIT-OFF 2-DOF	1.3 1.2 1.5	1.2 1.2 1.9	1.6 1.4 *	2.0 1.6 2.1	1.0 0.8 1.1	1.1 0.9 1.5	1.2 1.1 *	2.1 1.9 2.0	E1 E1 E2
SAMPLER 3-DOF	1.7 1.7 1.6	1.7 1.7 0.6	1.2 1.2 1.2	1.8 1.6 1.7	1.8 1.7 0.7	0.8 0.7 0.7	1.3 1.3 1.3	1.9 1.7 1.8	E2 E3 E1
SCOPPER 2-DOF	0.5	0.6	1.9	4.8	0.5	0.6	1.9	4.9	E1
RETURN-TOOL 6-DOF	1.9 1.1 1.9	1.0 1.5 2.6	3.2 2.8 *	3.4 12.0 16.9	1.9 1.9 1.7	1.0 1.1 2.4	3.0 2.3 *	10.6 11.3 16.3	E4 E1 E1
GET-TOOL 6-DOF	1.9 1.0 1.3	0.9 1.1 2.6	3.2 2.3 *	17.3 12.5 15.3	1.9 1.1 1.4	1.2 1.1 1.7	3.2 2.4 *	11.7 12.4 16.0	E4 E1 E1

master-slave with force feedback rarely benefits from supervisory control, (2) master-slave without force feedback can profit from supervisory control in tasks which require force feedback, and (3) both forms of rate control can be aided by supervisory routines regardless of the task.

In all cases the error rates for supervisory control were less than manual control. However, an interesting error was noted during the sampling experiments - occasionally the subjects pressed an incorrect button sending the sample to the wrong bucket.

Theoretically there is no reason why master-slave with force feedback should be any faster than supervisory control. Consider that the computer could simply mimic the human operator's best trajectory, and hence, be at least as fast. Unfortunately, in practice there is always a certain overhead associated with retransformation of coordinates, trajectory calculations and sensor logic. Also, it was generally observed that the subjects were making adaptive, orchestrated motions, whereas the computer was limited to more rigidly defined trajectories and states. In light of these observations it can be said that the faster master-slave times make more of a statement about the direction that future studies dealing with supervisory control should take than they do about its potential in teleoperator systems.

Although the experiments were not designed to measure the effectiveness of supervisory control during extended periods of manipulation, an interesting observation was made after the experiments had been completed - the manual experiments had been performed with rest periods between each run because the subjects complained of fatigue and boredom, while the supervisory experiments had been unintentionally run back-to-back since fatigue and boredom were not noted. From these observations it could be surmised that as a task becomes more involved and complex, boredom and fatigue will become increasingly important factors, tipping the scales even further in favor of supervisory control. However, experiments to validate this statement have yet to be performed.

9. Conclusions from SUPERMAN Experiments

Even under "ideal" control conditions supervisory control was found to be more efficient and effective (as determined from the task completion times and manipulation errors) than switch rate control, joystick rate control, and master-slave position control. Bilateral force-reflecting master-slave was found to be slightly faster than supervisory control, but more prone to errors. Since the experiments were performed under "ideal" conditions, it can be reasonably predicted that supervisory control will show even more advantage when used with degraded sensor or control loops (e.g., time delays, limited bandwidth, etc.), though the latter experiments remain to be done. In addition, an a posteriori observation of the experimental procedure appears to indicate that the effects of operator fatigue and boredom during extended periods of manipulation can be significantly reduced through supervisory control.

B. An Improved Supervisory Command System: MMIT^{*}

This section is based on Yoerger and Sheridan [1983] and Yoerger [1982]

1. Introduction

Following Brooks' initial work we built a much refined system for supervisory control of manipulators for underwater vehicles, primarily in the context of close-up inspection. The successor system we call MMIT for Man-Machine Interface for Telemanipulation. It consists of a movement control language and a computer graphic interface.

Using the MMIT system, an operator can teach the computer how to execute tasks after which the operator then monitors the system's progress. One of the most important features of the interface is that it can be used by operators with no physical understanding of manipulator control. The MMIT system was not designed to act independently of a human operator or as an automaton, rather it is a method for extending manual control.

This section provides a description of the interface design. It also describes an experiment performed to test MMIT in a simulated remote inspection task. The results from that experiment show great promise for practical applications of the system. Three general results are particularly important:

1. Task performance in terms of accuracy was improved under MMIT control.
2. The interface decreased the operator's dependence on visual feedback.
3. The system decreased the variability in performance between operators.

The MMIT system is currently in use at MIT and Naval Ocean Systems Center, San Diego. It will be used on RECON-V, a submersible being tested at MIT, and is currently a candidate for several other remotely operated vehicles.

2. System Design

There are several important features in the MMIT design. These include an advanced programming system that features analogic teaching, the combining of analogic data with symbolic commands, and the use of a dynamic simulation for visual feedback to the operator.

The programming system used in the interface allows the operator to teach the computer how to perform remote tasks. Tasks may be either preprogrammed, or new programs may be created during actual remote operation. The computer graphic display can be used to help the operator understand elements of the programming system without requiring a formal mathematical description of how the commands work.

The programming system is based on cartesian coordinate frames, as described for teleoperator tasks by Brooks. Absolute coordinate frames

(defined relative to the manipulator base) may be entered analogically, by pointing out their location with the manipulator. Relative coordinate frames (to be applied to an arbitrary coordinate frame) may be defined either analogically, or symbolically by entering a semantic description.

The analogic teaching capability is important in the system because it allows the operator to establish relevant positions and orientations quickly, transparently, and with minimal sensing on the part of the remote system. The operator moves the manipulator to the desired position and specifies a name for the position. A cartesian coordinate frame is then computed based on the current angles of the manipulator joints and is stored under the give name.

The arm may be repositioned at any of the defined coordinate frames (called a POSITION), or made to traverse along a series of such frames (called a PATH). This capability is useful for defining positions or paths that the arm should return to several times, such as the grasping position of a tool.

Analogic teaching can be especially useful when an operator does not know the exact coordinate values of a position, but can see via the graphic or television display what he wants. Using analogic definitions in combination with symbolic commands simplifies the teaching of teleoperation tasks for the operator. Programs can be written in which the operator points out an object in the task environment, and then describes an operation on the object by specifying motions built on defined positions.

Using the relative commands, the arm may be moved relative to its current position and orientation. These commands are useful for describing tool motions, such as turning a valve or brushing a weld.

Relative motions defined symbolically and absolute coordinate frames defined analogically may be used together very effectively. Using the analogic capabilities, the operator may define an object of interest as a series of coordinate frames. These absolute coordinate frames can then be used as reference frames for relative motions which define a task. An example, inspection of a weld, will be described later.

The computer can be instructed to make decisions during the control process using structured flow of control statements. Complicated tasks for the manipulator can be composed as a hierarchy of subtasks (Figure 8) through the extensibility mechanism of the FORTH language in which the system is implemented. Once the motions required for a particular task have been defined in the program and analogic data has been input regarding the environment, entire tasks can be accomplished with a single command.

A key element in the effective operation of the system is use of the dynamic computer graphic display, designed by Winey and developed further by Fyler (See Sections III-B and III-D, respectively). This display allows the operator to view an image of the manipulator and task environment from any angle. The operator can rotate, translate, or zoom the display. The dynamic display serves three major functions within this system.

(1). The display is a very effective aid for monitoring the actions of the manipulator and the condition of its environment. Because the images are simulated, the quality of the display is not dependent on the quality of a

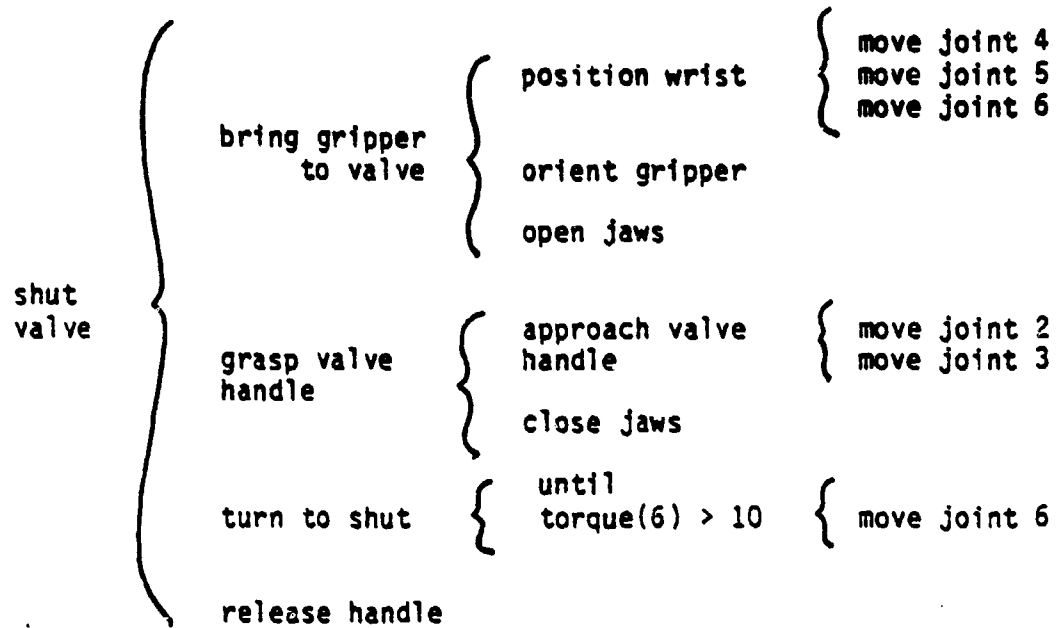


Figure 8 Warnier-Orr Diagram of Shut Valve task. This is similar to the procedural net model representation, but decisions may also included. In this example, the valve will continue to be turned until a preset value of torque is reached on the last manipulator joint

television picture.

(2). Motions of the manipulator can be simulated to test programs before they are actually executed. Such simulations can be run in faster than real time.

(3). The display can be used to show the results of computations done by the computer after high level descriptions have been given, thus helping the operator understand how the system works. Defined coordinate frames corresponding to key parts of the environment may also be displayed.

It is important to note that the operator remains in control of the system during the execution of all tasks. If he detects a problem of which the system is unaware, he can interrupt and assume control manually or he can invoke computerized functions to get out of trouble, such as repeating the last several computer-controlled moves in reverse order.

3. Experimental Evaluation of MMIT

Experimental Design - An experiment was designed to evaluate a working model of the MMIT system. The experiment explored how well an underwater inspection task could be performed with a master-slave manipulator under three different methods of control. Performance was compared for one manual control mode and two supervisory configurations.

The experimental task was the remote inspection of an underwater weld. This task was chosen because it has proven difficult to perform under manual control in actual underwater tests and includes many representative elements of supervisory control which can be measured and perhaps generalized. Remote weld inspection usually involves two steps. The weld is first cleaned down to bright metal and then some form of non-destructive testing (NDT) is performed, the most common being still photography. Both steps share the same basic requirements - that a tool (cleaning jet or camera) be moved through a trajectory that traverses the weld, that keeps the tool pointed at the weld, and that maintains a specified distance between the tool and the weld.

Performance of the task was evaluated qualitatively in terms of completion time and the error of the manipulator arm in following the desired trajectory.

Three control modes were evaluated in the experiment. The first was standard master-slave with force feedback. The second and third modes supplemented the master-slave control with different types of supervisory computational aids.

Mode 1: Master-slave control. This mode is generally considered to be the mode with the best performance of any purely manual control mode. The computer was not used for control.

Mode 2: Analogic teaching. In this mode, the operator teaches the trajectory to be followed as a series of discrete positions using the master-slave control. The system can then move the arm through a smooth path which intersects all taught positions. The system does not transform the taught positions, but only interpolates between them, so the computational requirements are quite low. A major advantage over mode 1 is that the taught path may be repeated as desired. This mode is similar to the method used to

program many industrial robots.

Mode 3: Combined analogic and symbolic teaching. In this mode, the operator first teaches the weld (rather than the trajectory) as a series of discrete positions. The operator then invokes a procedure that symbolically describes a trajectory relative to the taught positions. The procedure attempts to create a trajectory that points at the weld while maintaining a specified distance between the manipulator and the weld. An advantage of this mode over mode 2 is that the trajectory may be repeated as desired with different distances or orientations between the analogic description of the weld and the generated trajectory.

During actual water jet cleaning underwater, the operator has difficulty seeing because the jetting action obscures his view. For this type of task, either of the supervisory modes (2 or 3) would have a distinct advantage over any manual control method by separating the task into teaching and execution phases. For both modes 2 and 3, the operator could do the teaching before the jet obscures his view, then the computer could carry out the jetting operation.

Experimental Setup - The experimental setup is shown in Figure 9. The manipulator used in the experiment was an E-2 master-slave arm. The computer was a PDP 11/34 running the MMIT software. The sensors and actuators of the manipulator were the same for all control modes.

A test weld was developed for the inspection task which included both straight and smoothly curved sections. Subjects viewed the weld and remote manipulator through a television system. The weld and camera were arranged to present a variety of viewing angles to the operator in order to test the sensitivity of the different modes to changing spatial relationships between the camera, manipulator, and weld. The specific task consisted of defining a trajectory which remained one inch from the weld and pointed directly at the weld.

Experimental Trials - Three experienced subjects were used in the experiment. Sessions for each subject consisted of three trials in each control mode with the test weld in three different orientations for a total of nine trials per session. Each subject performed three sessions, and data was recorded on the last session.

Performance Criteria - Accuracy was evaluated by two criteria as shown in Figure 10:

1. The shortest distance between the tip of the manipulator tool and the weld was computed as a function of time, with RMS computed for each run. For perfect performance, this distance would have been maintained at one inch.
2. As a measure of orientation accuracy, the distance between a line oriented with the hand and the closest point on the weld was computed. This distance corresponds to the distance the center of water jet would miss the center of the weld. RMS was computed for each run.

For these measures, estimates of the measurement noise were obtained. Scores for all modes were substantially higher than the measurement noise.

Experimental Results - Analysis of variance showed the effect of control mode was significant for both the position and orientation criteria

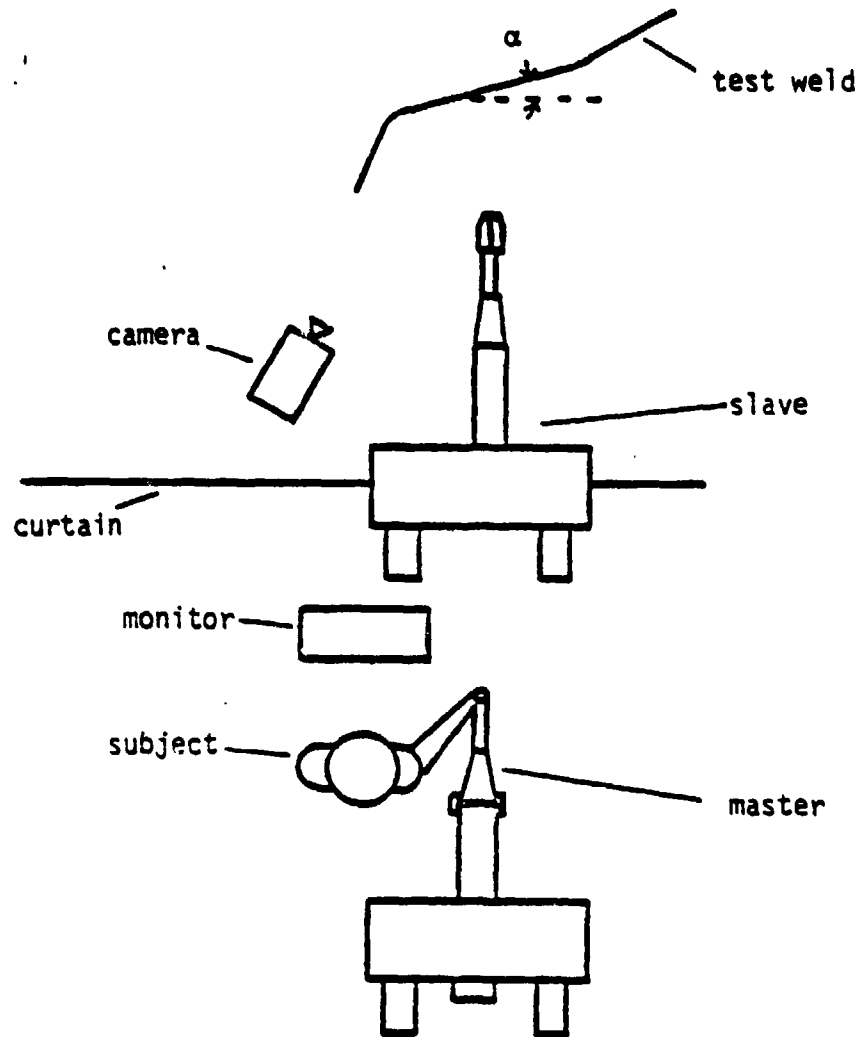


Figure 9 Experimental setup. Three orientations of the test weld were used ($\alpha = -15^\circ, 0^\circ$, and 15°).

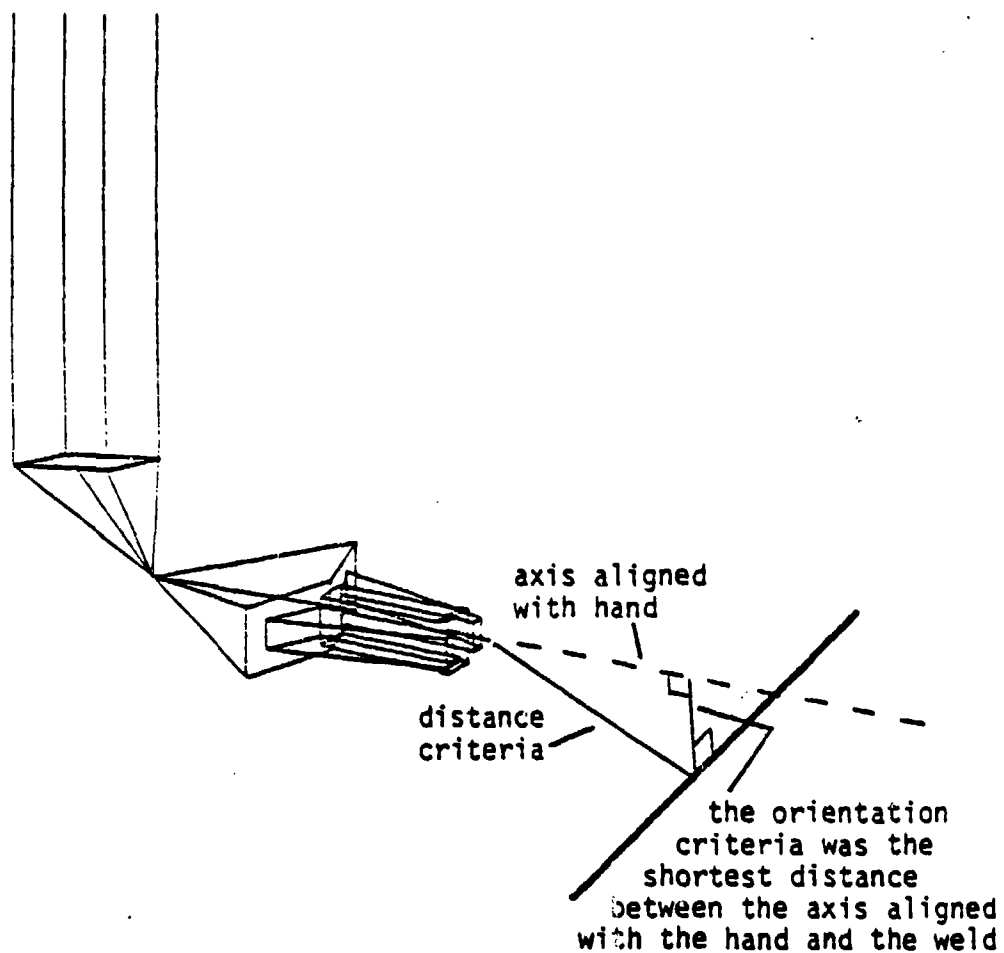


Figure 10 Performance criteria. The distance criteria was the shortest distance between the tip of the tool and the weld bead. The orientation criteria was the shortest distance between the axis aligned with the hand and the weld bead.

($p < 0.025$). Examination of plots of subject means and RMS errors for both position and orientation criteria (Figures 11 and 12) show that each subject improved similarly across the control modes. In general, the largest improvement was between modes 2 and 3. Mode 3 also showed the lowest variation between subjects.

Figure 13 shows the performance criteria as a function of distance along the weld, averaged across all subjects and runs. Large systematic changes in the errors can be seen for modes 1 and 2 as the spatial relationship between the camera and the weld changes. For mode 3, performance was much more uniform despite the large changes in the quality of the visual feedback.

4. Conclusions

This section has described how the MMIT system works and reports on an experiment that shows the system's usefulness in improving performance in a simulated remote inspection task. The particular task was patterned after cleaning and inspecting a curved weld. The system demonstrated improved performance in terms of accuracy, decreased dependence on the quality of visual feedback to the human operator, and decreased variability between individual operators over more conventional approaches.

The computational and sensing requirements of the system are quite low. The system requires standard 16 bit microcomputers (LSI 11/23 for example) and needs only joint position sensing of the manipulator, making it practical for implementation on remotely operated vehicles. As much as possible, the system was designed to be manipulator independent.

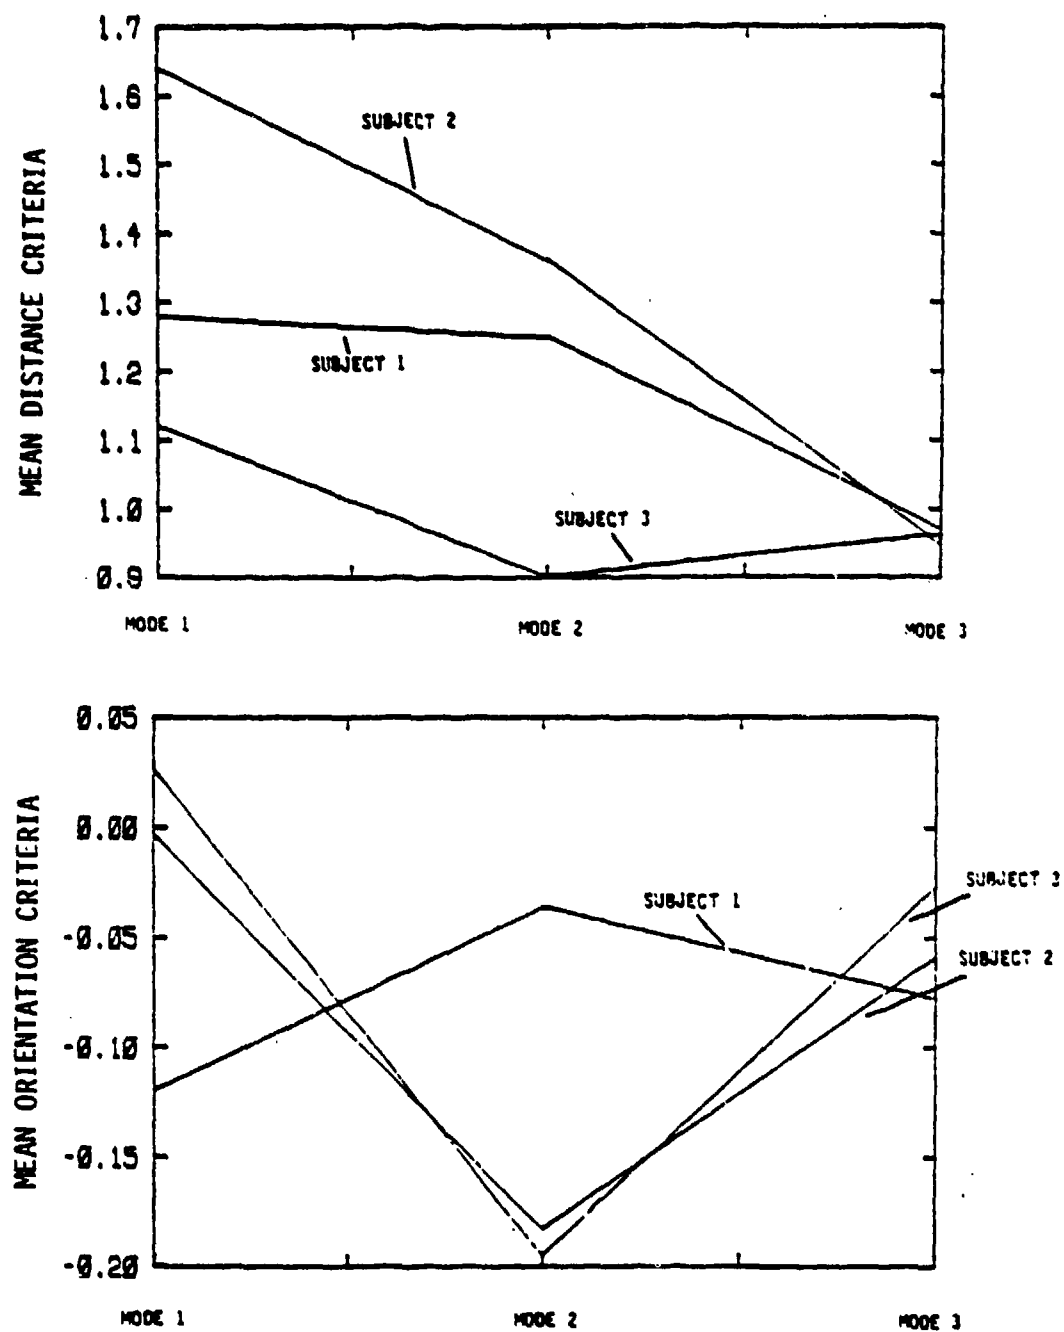


Figure 11 Mean performance as a function of control mode for the distance and orientation measures. No significant effects were seen.

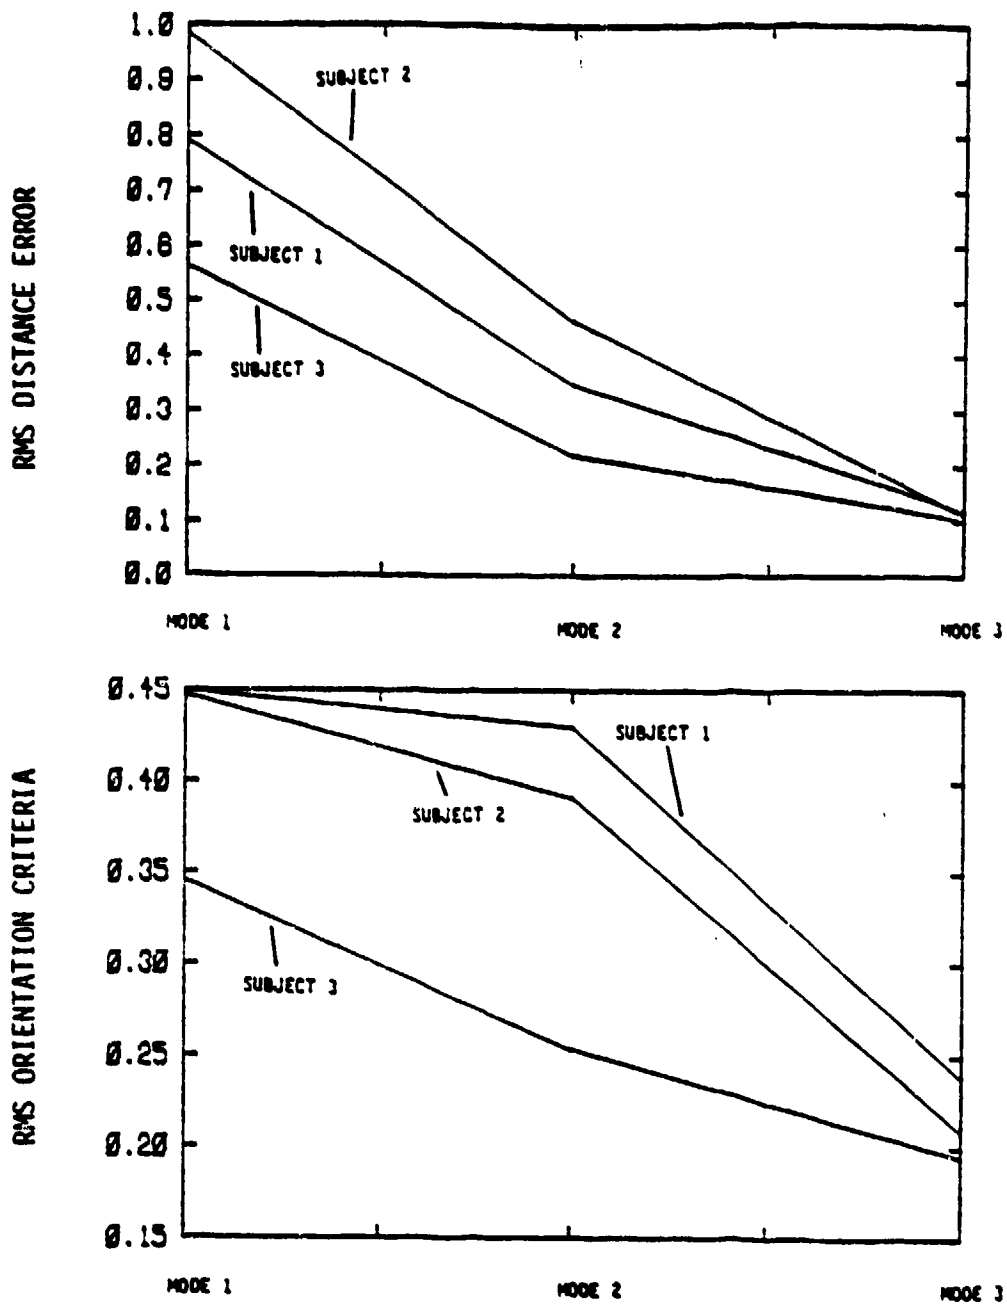


Figure 12 RMS performance as a function of control mode for the distance and orientation measures. The effect of control mode was significant for both measures. The variation between subjects was much less in mode 3 for both measures.

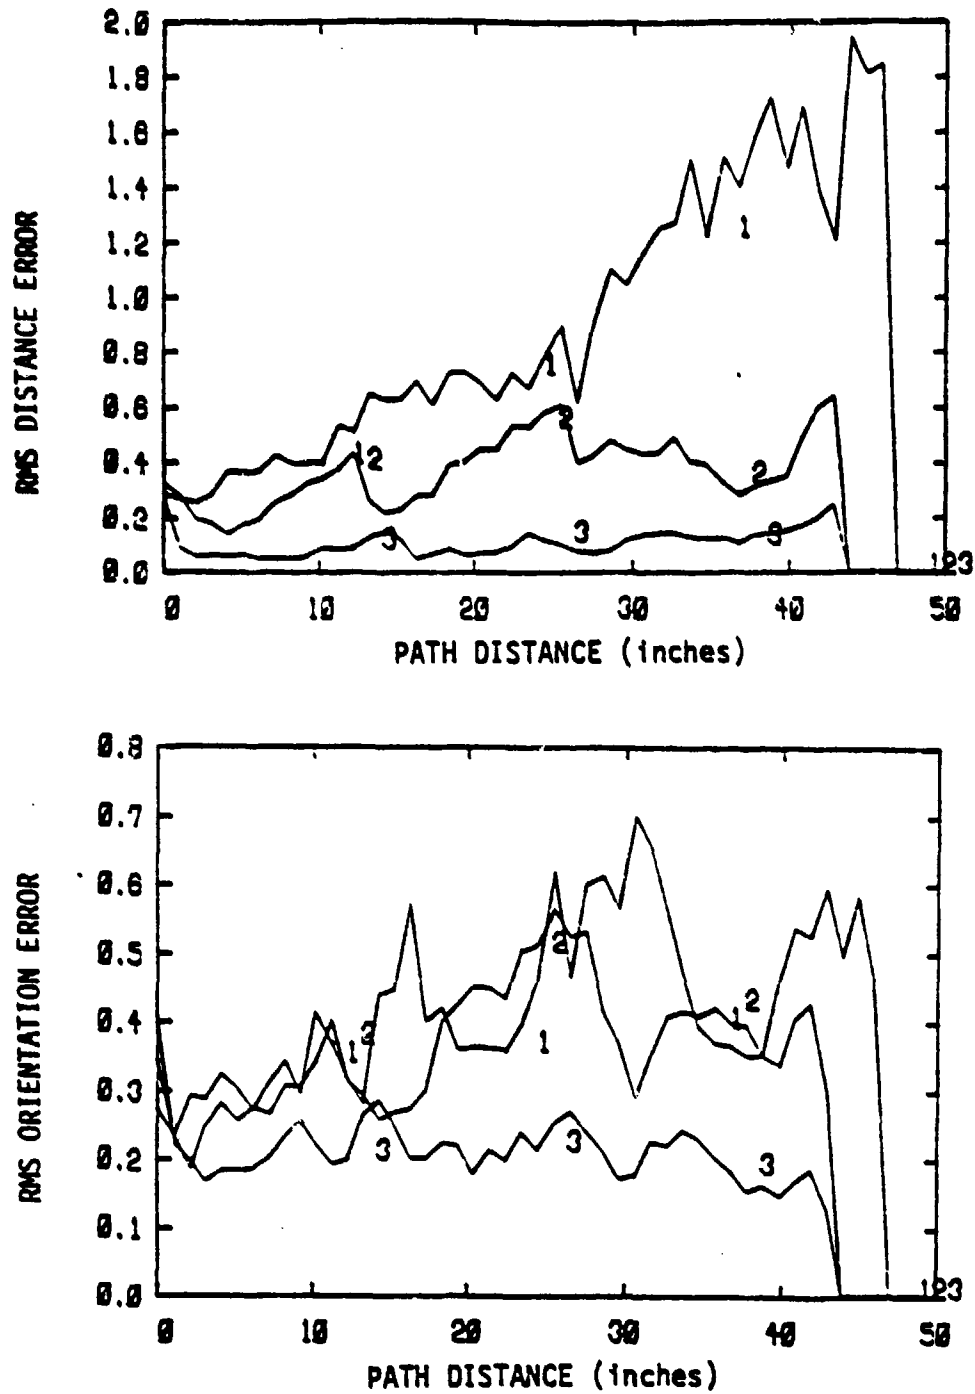


Figure 13 Performance as a function of distance along the weld. The upper plot is for the distance measure, and the lower plot is for the orientation measure. On each plot, curves for each mode are shown.

C. Factors Affecting Pointing: Telling The Computer About Geometry

* This section draws upon material from Yoerger [1982] and Tzelgov, Yoerger and Sheridan [1983]

1. Introduction

This work looks at the factors affecting the accuracy of "pointing" toward some prespecified direction in space. To be more specific, accuracy in orienting the manipulator's arm normally or at a specific angle to some plane at some prespecified location are important for some teleoperation tasks. Both Brooks' [1979] and Yoerger's [1982] work showed that a complex trajectory may be defined by a computer if the cartesian coordinates of a number of anchoring points along the trajectory can be provided by the operator. To provide these coordinates, the operator has to locate the manipulator arm at each point in a prespecified (often normal) orientation. Thus, in this specific application, angular accuracy is of critical importance. In these experiments we investigated the operator's ability to orient the manipulator's arms normally to a plane at prespecified points differing in their spatial locations.

2. Experiment 1: Direct Viewing, Varying Task Orientation

This experiment focused upon the ability to orient the manipulator's arm under direct viewing conditions. It is well known that manual motions performed on the mesial plane are more accurate than sidewise motions. To test if a similar relationship holds for angular accuracy, performance while orienting the manipulator normally to points on the mesial plane as well as on side planes was evaluated. We assumed that accuracy in different tasks should be rather insensitive to relative displacement of the viewpoint along the vertical axis. To test for this effect we used several planes differing in their inclination. And finally for each combination of horizontal displacements and inclinations five points differing in the relative x,y coordinates were evaluated.

Experimental Design- The Argonne E-2 master-slave force-feedback manipulator connected to a PDP 11/34 was used for the experiments. The experimental task was to orient the slave manipulator hand normal to a plane at a specified location on the plane. Three different planes were used. Each was at a different position, Figure 14, and each was inclined by a different angle, as shown in Figure 15. Each position was equidistant from the base of the slave manipulator. On each plane, five locations were marked. Inclination of the plane, lateral position of the plane, and locations on the plane were taken to be the independent variables.

A 3x3x5 full factorial within-subjects design was used. For this type of design, each subject performs the task for each combination of the independent variables. The within-subjects design was chosen because it allows effects to be observed for a small number of subjects despite interaction between subjects and main effects.

Three subjects were tested, all right handed male engineering students with normal or corrected vision. Each subject performed three repetitions at

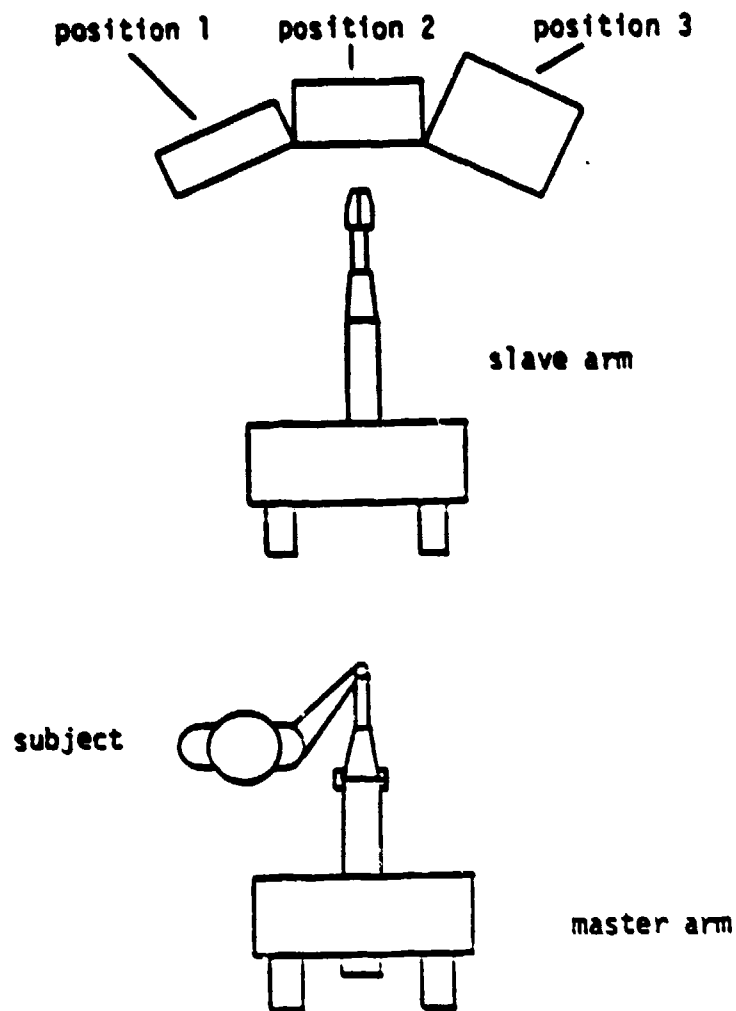


Figure 14 The direct viewing experiment: each plane could be placed in any of three positions. Each position was equidistant from the manipulator base. The center position was directly in front of the

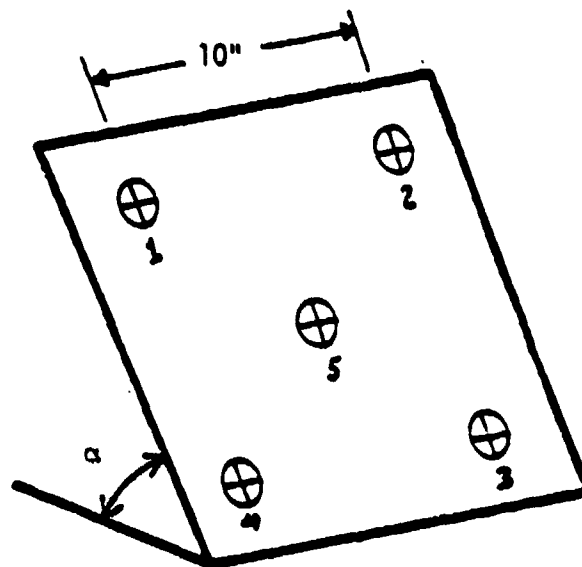


Figure 15

Three different planes were used in the experiment. Each was inclined at a different angle α to the horizontal plane. Values of α used were 30, 45, and 60 degrees. On each plane, five locations were marked.

each combination of inclination, position, and location, so that each subject performed a total of 135 trials in a randomized order. The trials for each subject were broken up into three blocks, with each block consisting of a fixed combination of inclinations and positions. The order of these blocks was counterbalanced across subjects.

Two components of angular error were analyzed which together describe the projection of a unit vector attached to the manipulator hand onto the task plane to be defined. Together, these two components give the magnitude and direction of misalignment. They will be called x error and y error. A perfect (perpendicular) alignment gives no projections. It is a point at the origin.

Procedure - Each subject was given the same written description of the task and how his performance would be judged. The instructions emphasized that accuracy was the prime performance measure, although performance time would also be recorded.

Each trial began with the master arm locked in computer control in the same position. The experimenter then told the subject at which plane and location on that plane the subject should position and orient the arm. The manipulator was then placed in manual control and the timer was started. The subject indicated when he had positioned and oriented the arm to his satisfaction by depressing a hand-held pushbutton switch. The timer was then stopped, and the current values of the manipulator's joint angles, the elapsed time, and the commanded position and location were recorded.

Results - For x error and y error, a method was devised for displaying both the mean and variance in a meaningful way. Examination of the data shows that the error does not vary independently in x and y. The variance of this two dimensional error is best described by a covariance matrix:

$$\underline{\Sigma} = \begin{vmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{yx} & \sigma_y^2 \end{vmatrix}$$

where the diagonal elements are the variances, and the off-diagonal elements are the covariances. In general, it is possible to find a set of coordinates for which the covariances are zero. The angle of the principal axes may be computed from the relationship:

$$\theta = \frac{1}{2} \tan^{-1} \left(\frac{2\sigma_{xy}}{\sigma_x^2 - \sigma_y^2} \right)$$

The values of the variances along these axes may be called the principal variances. The principal variances are uncorrelated measures of the spread of the data. The values of the principal variances may be found by the relations:

$$\begin{vmatrix} \sigma_x^2 \\ \sigma_y^2 \end{vmatrix}' = \begin{vmatrix} \cos\theta & \sin\theta \\ \sin\theta & \cos\theta \end{vmatrix} \begin{vmatrix} \sigma_x^2 \\ \sigma_y^2 \end{vmatrix}$$

The means and variability of the error data may be summarized by plotting an ellipse centered at the mean value of x and y error, with the major and minor axes of the ellipse equal to the square root of the principal variances (principal standard deviations). This plot provides a descriptive "error

footprint" for the data x-y. (Figures 16 and 17).

Analysis of variance for x error showed significant effects for both position ($F(2,4) = 19.9$, $p < .01$) and location on plane ($F(4,8) = 11.9$, $p < .01$) while the effect of orientation was not significant.

A Newman-Keuls post-hoc test was used to test for significant differences between individual position means. This test showed that the mean x error was significantly different for each of the three positions.

The effect of position was significant for some y errors ($F(2,4) = 8.09$, $p < .05$). For y error, there was no significant difference between positions 1 and 3 (the left and right positions). Both positions 1 and 3 differed significantly from position 2 (the center position).

Interpretation of the error data for the position main effect can shed light on the relative importance of perceptual and motor considerations. The left and right positions, positions 1 and 3, are similar from a perceptual point of view, as the operator was positioned directly between these two positions. Position 2, the center position, was directly in front of the subjects, quite different, as the master slave manipulator is a right-handed device. All three means are shifted to the left due to common motor considerations.

3. Experiment 2: Television Viewing Experiment, Varying Camera Viewpoint

The previous experiment looked at the effect of position and orientation of the defined frame for a generally fixed viewpoint. In this experiment, one orientation was tested for different viewing angles. The task and experimental setup was the same as in the direct viewing experiment, but television viewing was used. Four different camera positions were used, as shown in Figure 18. Only one plane in a fixed position was used, as the previous experiment indicated that motor considerations are less important than perceptual factors. Again there were four locations marked on the plane. Camera position, location on the plane, and the practice effect were the independent variables.

Experimental Design - A full within-subjects factorial design was used. Eight subjects were tested. Each subject performed 8 blocks of 16 trials. Within each block, the subjects made 4 trials for each location on the plane in a different randomized order. After each block, the camera position was changed. The first four blocks corresponded to the first phase. In the first phase, the subject was given feedback about his performance from the graphic display after performing the trials for each block. In the second phase, the subjects performed blocks for each camera position again, but without feedback from the error display. Within each phase, the order of camera positions were counterbalanced across subjects. Each subject was a right handed engineering student with normal or corrected vision. The dependent variables were elapsed time and orientation errors as defined in the previous experiment.

Procedure - Each subject was first given written instructions, emphasizing the same performance measures as in the direct viewing experiment. Subjects were also given written instructions about the meaning of the error

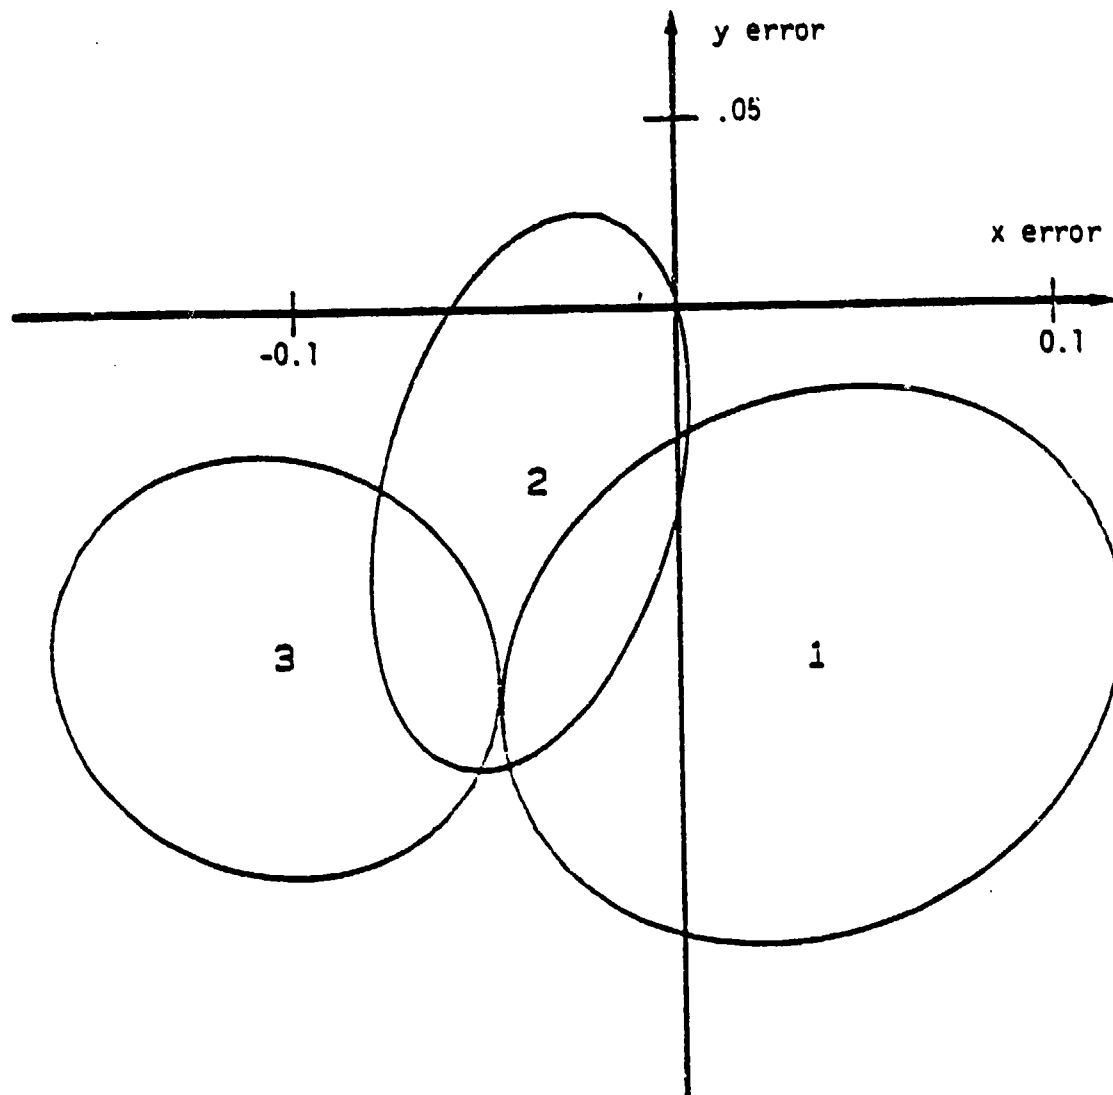


Figure 16 Orientation errors by position of plane. Positions 1 and 3 showed significantly higher y and radial error than position 2. For x error, all positions differed significantly.

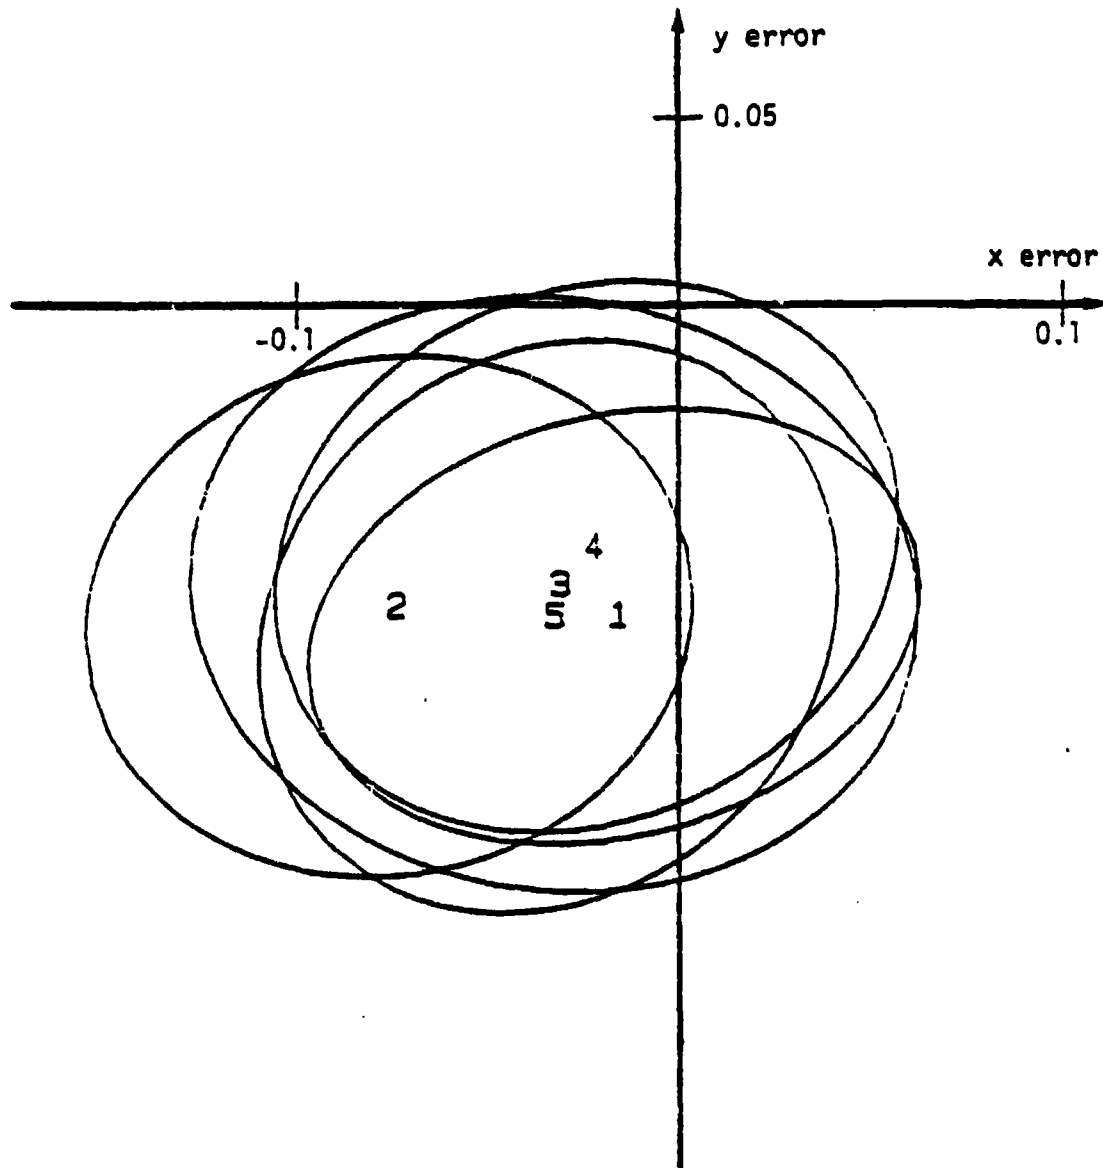


Figure 17 Errors by location on plane. No significant effect was seen.

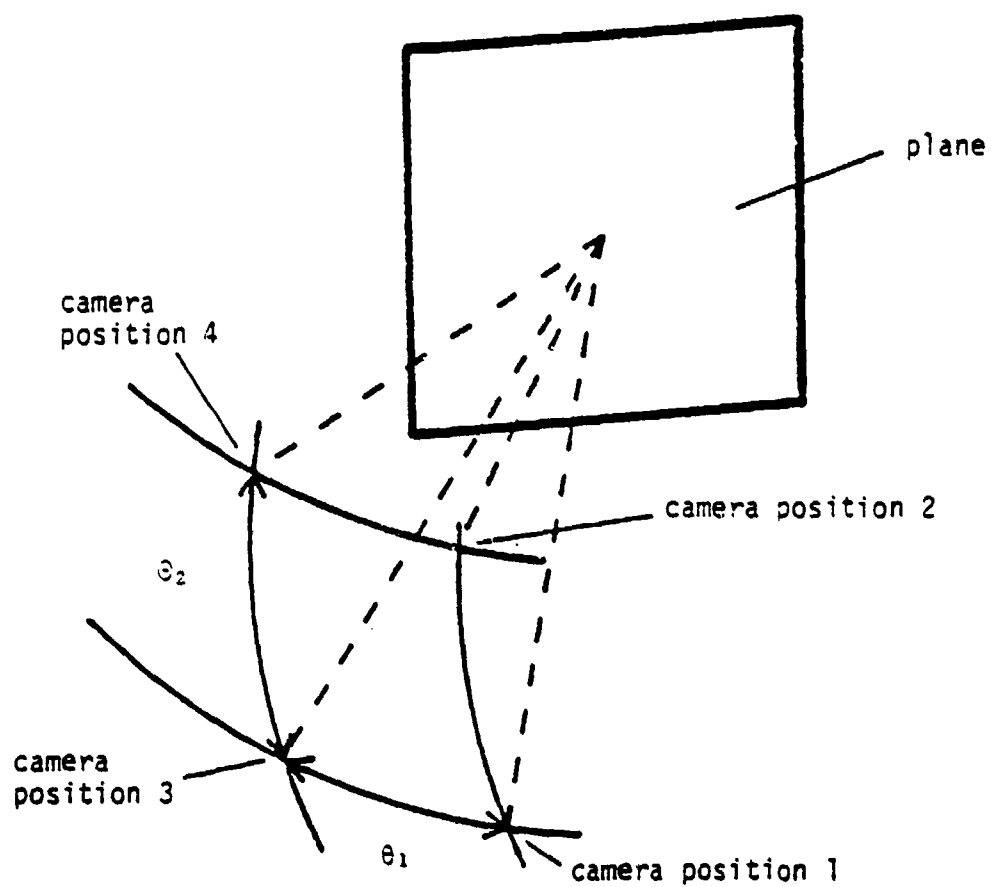


Figure 18 Camera positions were defined in a spherical coordinate system based at the center of the plane to be defined.

display. Each trial proceeded in the same manner as in the earlier experiment.

Results - Analysis of variance was performed on both x,y, and radial error and time data. Error plots were also produced, as was done in the earlier experiment.

Figure 19 shows the error data for the different camera positions. The difference was significant for y error ($F(3,21) = 41.$, $p < 0.001$). This difference in y error is shown more clearly in Figure 20.

Figure 20 shows the error plots for the different locations on the plane. This difference was significant y error ($F(3,21) = 50.$, $p < .001$). Means for each combination of camera position and location on the plane were also tested. This interaction was barely statistically significant. Radial error, i.e., combined magnitude of the angular error, showed only small changes for each combination of camera position and location on plane.

Figure 21 compares the TV and direct viewing experiments. The mean and variation averaged across all independent variables and subjects for TV viewing is shown, along with the mean and variation for similar inclination of plane, and locations on plane for direct viewing.

The direct viewing experiment showed that this task is dominated by perceptual considerations. The television viewing experiment provides further clues about the source of these errors.

The significant differences in performance as a function of camera position confirm the perceptual nature of this task. For each location (i.e. constant motor factors) y error varied significantly. The small interaction between location on plane and camera position shows that the location effect was fairly consistent for different camera positions.

The significant differences in performance for different locations on the plane can also be tied to perceptual effects and seem to support the idea that the direction of the error is determined by the viewpoint, as seen in the direct viewing experiment. The mean error for different locations on plane (Figure 20) is directly related to the actual positions on the plane. The errors for location 1 (upper left hand corner of the plane) are consistently toward the lower right hand corner, and so forth.

Examination of the plot of y error (Figure 22) shows several effects. Raising the camera always made y error more positive, while reaching for a lower location on the plane also made the y error more positive. This is the same bias toward the viewpoint which was seen in the direct viewing experiment. This corresponds to a consistent underestimation of the orientation of the plane relative to the direction of gaze, even when the direction of gaze is defined by a camera. This underestimation effect is consistent with the underestimation of radial direction constancy found by Hill [1976].

4. Conclusions from Direct Viewing and Television Experiments

The following conclusions can be drawn from these experiments:

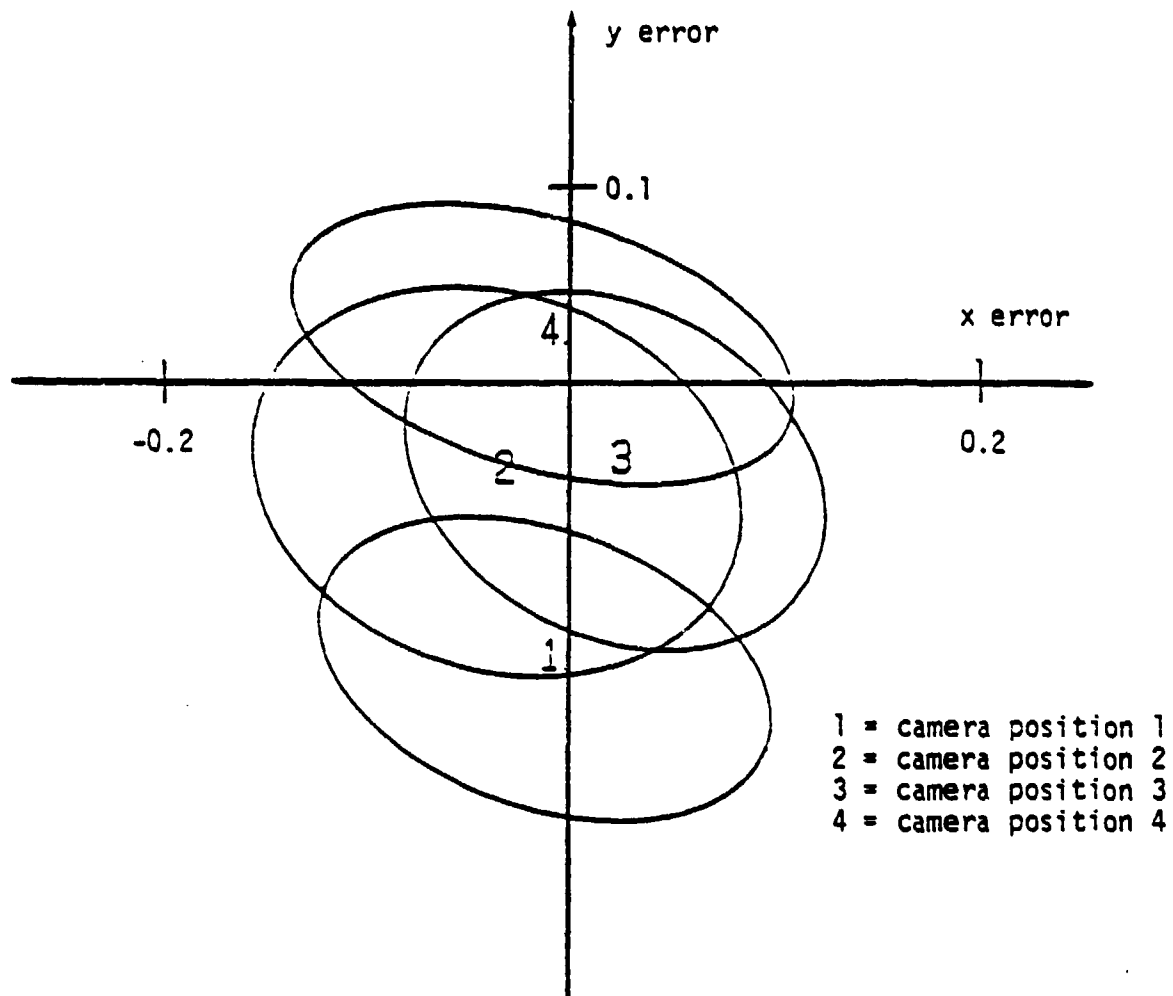
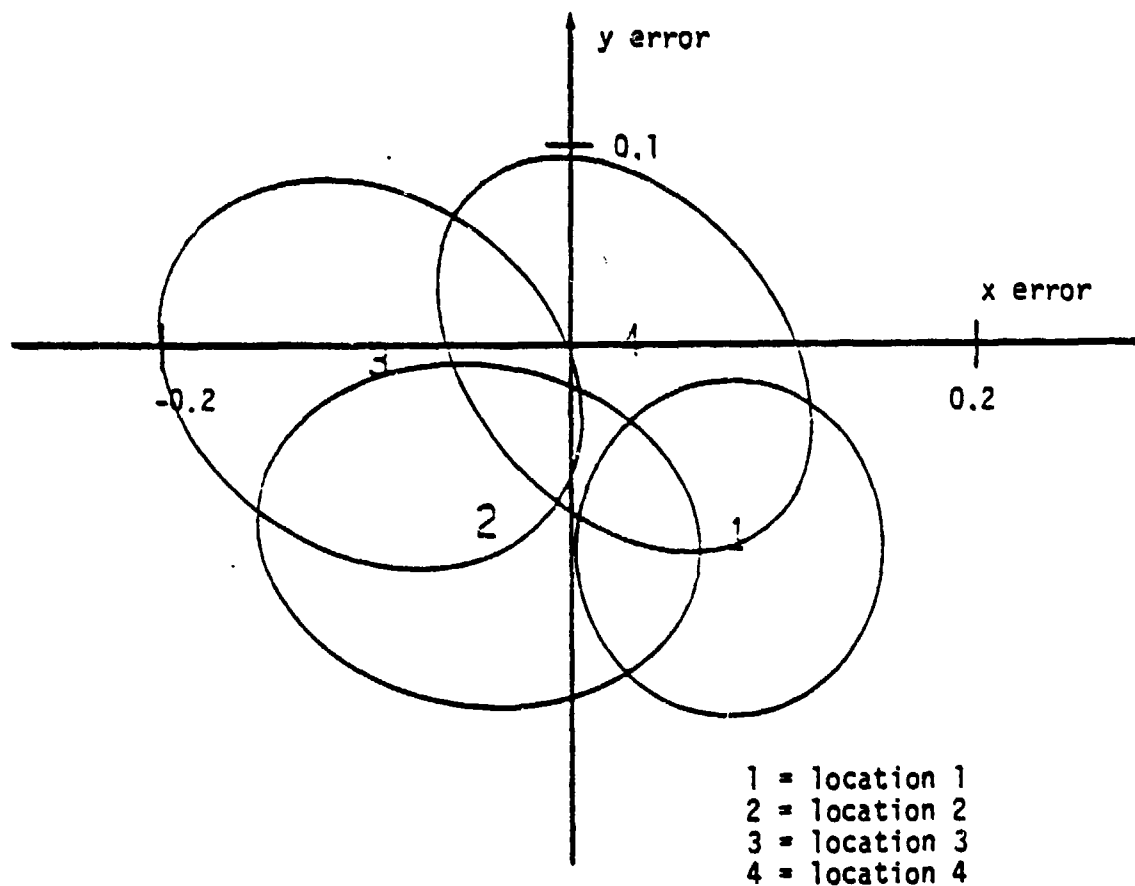


Figure 19 Error plots for different camera positions. The effect was significant for y error, but not for x error.



locations on the plane
were defined in figure 6.2

Figure 20 The effect of different locations on the plane can be seen in this plot. The effect of location on the plane on x, y, and radial error was significant.

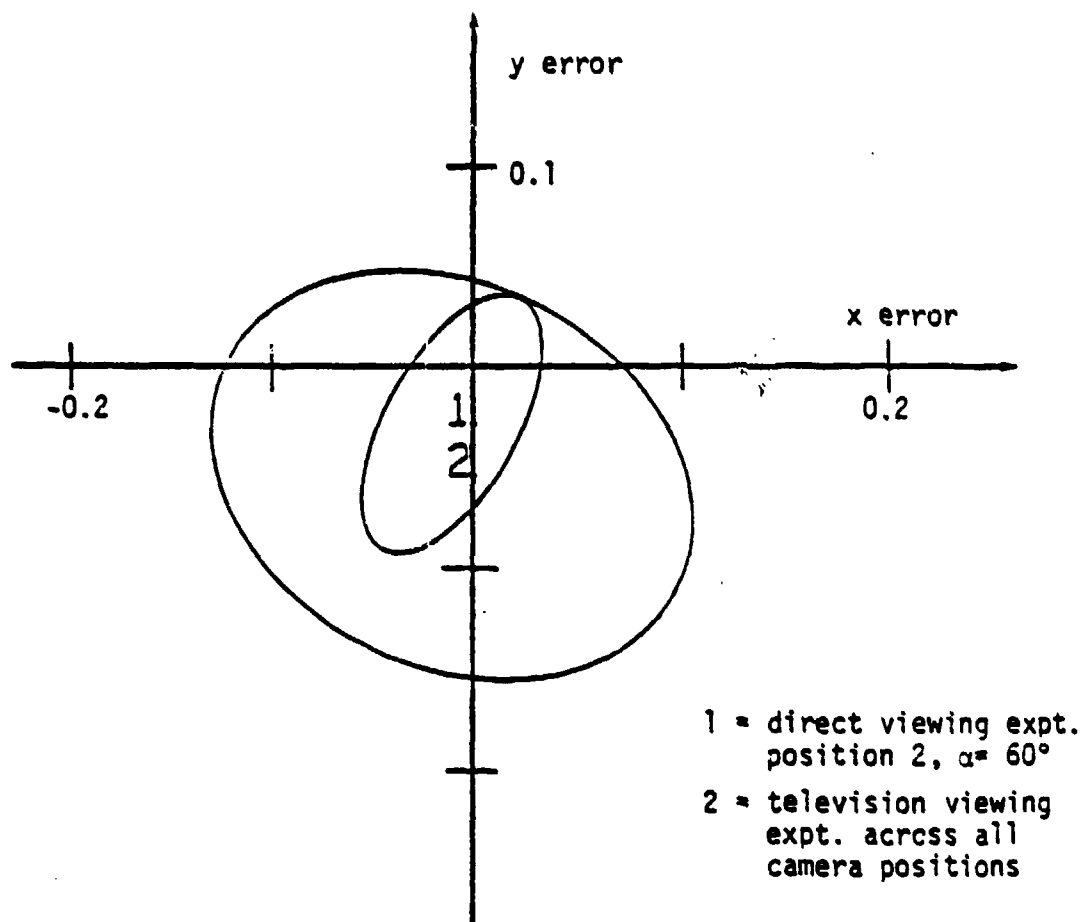


Figure 21 Comparison of direct viewing and television viewing experiment for similar inclination of plane, position of plane, and location on plane.

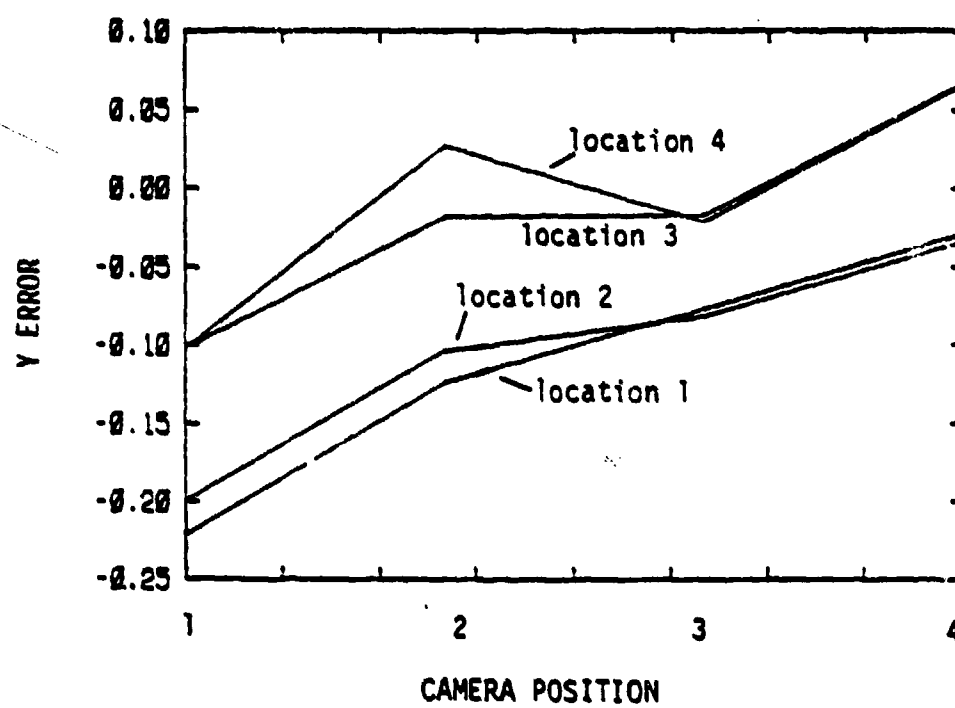


Figure 22 Y error for different camera positions.

1. While the average error with TV viewing was similar to the direct viewing experiment, the variation in error was larger.
2. The direction of error was consistently tied to perceptual issues corresponding to the spatial relationship between the plane to be defined and the direction of gaze.
3. For both direct and television viewing, subjects consistently underestimated the relative orientation between the direction of gaze and the plane to be defined.
4. A 45 degree angle between the direction of gaze and the plane to be defined was found to have best performance.

D. Automatic Compensation For Motion Disturbances In Teleoperation*

* This work is based on Tani [1980] and Hirabayashi [1981]

1. Introduction

This short section is included to give an example of a fairly straightforward form of partial automation or aiding which has been implemented in aircraft and spacecraft, power plants and other forms of human-supervised, semi-automatic systems. These are situations where relatively straightforward and conventional feedback control is employed to minimise disturbances which occur with respect to some variable while that same variable is being controlled in parallel or at a higher level by a human supervisor or machine.

Our example in this case occurs in remote manipulation of objects undersea where the manipulator base is sometimes a mobile submarine or vehicle which may move relative to the object being manipulated, and this makes either direct manual control or supervisory control difficult. This relative motion occurs either because a manipulator is being supported by a vehicle which is hard to hold steady against ocean currents or other disturbances, or because the object being manipulated is being buffeted, or both. The same problem could occur in space or in terrestrial mechanical manipulation.

A means to overcome this is to make some measurement of the relative changes in displacement and orientation between manipulator base and object, either by optical, sonic or mechanical means, then to compensate for these changes by added motion of the end effector. The use of a mechanical "measurement arm" is one approach. Other means are optical, sonic, etc.

2. Experiment.

Hirabayashi [1981] implemented a measurement arm compensation scheme experimentally. He constructed a six degree-of-freedom (all angular movement) measurement arm which was lightweight and flaccid (offered little restraint). A six-degree-of-freedom Jacobian matrix transformation then allowed determination of the relative displacement of any object to which the measurement arm was attached.

Using a task-board with holes into which pegs were to be inserted, Hirabayashi drove the task board with a continuous random positioning device (three degrees of freedom, roughly 0.2 hz bandwidth, 6 inches root-mean-square amplitude.) He then attached the measurement arm to this task board, and used the resulting measurement of displacement to produce a compensatory displacement bias between the master and slave.

When the arm was under computer control it compensated to within 0.2 inches, even with a crude three-foot-long measurement arm. Then computer compensation was added to manual master-slave control of actions relative to the moving object. It was found to be much easier with the compensation than without it to put pegs into the holes in the moving task board.

E. Allocating Machines Versus Doing It Yourself*

* This section is based on Wood and Sheridan [1983] and Wood [1982]

1. Introduction

This final set of experiments was relatively abstract in nature. We were interested in cases of man-machine systems where the human operator is given responsibility for optimizing system performance by combined allocation of both his own time and the time of various machine aids subordinate to him. All the tasks can be done manually if necessary, but time is constrained. Use of the machine aids may well enhance productivity. We considered only machine aids which could, once a task was assigned, complete it on their own. Eliminated, therefore, were machine aids like lawn mowers or pencil sharpeners which require continuous human control.

In the specific experimental systems examined here a human subject (operator) is faced with a variety of tasks to accomplish (or task opportunities) presented to him on a computer graphic display. He is given a number of machine aids to supervise. That is, he can assign the aids to do the tasks or he can do them himself, in either case by pressing appropriate buttons on a key pad.

It was deemed reasonable to believe that as the operator makes decisions and assigns tasks he seeks to maximize some reward function. For simplicity, a unidimensional reward function was assumed. It can be argued that in real life, people seek to maximize a variety of attributes like money, happiness, and respect. However, on a task basis it seemed an acceptable approximation to say that tactical decisions are based on maximizing a single objective.

Having hypothesized a unidimensional reward it is simplest to assume a linear utility for such reward. This assumption is crucial to allow exploration of the cognitive interface. In the experimental situation sub-optimal performance by the subject should reflect some psychological barrier preventing the operator from fully grasping the complexities of the task environment. However, without a reasonably linear utility for the reward, it may be that the operator is trying to maximize a different reward function than that of the supposed optimal model. That is, the experimenter's optimal and the subject's optimal would not be the same.

2. Preliminary Experiment - BOXCLR

We assumed that the reward in our experiments should be "points", and that linear utility would be a reasonable assumption. A simple experiment performed at a computer-graphic terminal was devised to test this assumption. In the experiment, named BOXCLR, the operator was made aware at the outset of all tasks to be accomplished, all rewards, and all costs. By eliminating the subject's uncertainty and assuming he still knew how to optimize, we assumed that any sub-optimal behavior was attributable to some non-linear utility for the reward, presumably because of some misconception about costs or rewards.

In BOXCLR, the subject was given a number of boxes or tasks to "clear" from the computer screen. This was done by assigning either a "human" (himself) or a "machine" (it looked like a bulldozer). The machine accomplished the task faster but cost more. Only one machine of each type could be operating at a time. When all boxes were cleared the experiment ended.

Three subjects each conducted 96 BOXCLR trials in which holding costs for a task not yet done were fixed, and where machine wages and the number of boxes varied. In this mini-experiment the subjects performed quite close to the optimal. A slight decline in performance was noted as machine wages increased. This may indicate that the operators did not have perfectly linear utility functions, but they were certainly close enough to justify use of the assumption in the experimental paradigm.

3. A Multi-Queue Allocation Experiment - SUPER

In this experimental paradigm subjects were faced with a dynamic multi-task environment where machine aids were available to assist the operator. To simulate the mental and physical separation of the operator from all but the task currently being performed, work areas were created and the subject could only look into one area at a time. This makes the experiment "multi-queue", with the operator searching from area to area for tasks. The experiment was performed as a simulation with information appearing before subjects on a computer screen "playing field". The computer display or diagram of the field is seen by the subject shown in Figure 23.

In the experimental game there are R classes of tasks that can possibly arise. For simplicity, only one member from each task class can appear at any time. Each task class arises in its specific work area, and only one area is displayed at a time. If a task exists in a work area, it is signified by the display of a box in the work area. To complete the task the operator must move the box to the right end of the area. The box may be moved manually by pressing the "DO TASK" button, or may be given to a machine aid by pushing the "ASSIGN MACHINE" button. The operator also has the option of changing work areas. This is accomplished by pressing the control button for the desired new area.

A subject on SUPER faces only two types of decisions. If a task is before him, should it be performed manually or by machine? And, if there is no task, should he leave this work area to look at another? In making these decisions the subject must take into account all the variables which currently define the state of the system.

Three independent ratios using these parameters were found to characterize each task. These are:

- the ratio of reward to holding cost,
- the ratio of service rate of tasks to the arrival rate of tasks
- and the ratio of transition time between work areas to the mean arrival time of tasks.

The best strategy for two cases with the same values for these three ratios will be the same regardless of the absolute values of the task state variables.

4. Determining an Optimal Model for SUPER

It is fairly simple to determine an optimal model for operator behavior in a situation with very few variables. However, with increasing numbers of task classes, machine productivity levels and work areas, directly calculating the value of each combination of actions becomes too complex and some simplification must be used.

A decision tree maps out all possible paths open to the operator at each decision point. The reward accumulated for each unit of time is the most efficient method of evaluating a strategy's effectiveness. In simplifying calculations we sought a system which would provide an expected reward per unit time (RPT).

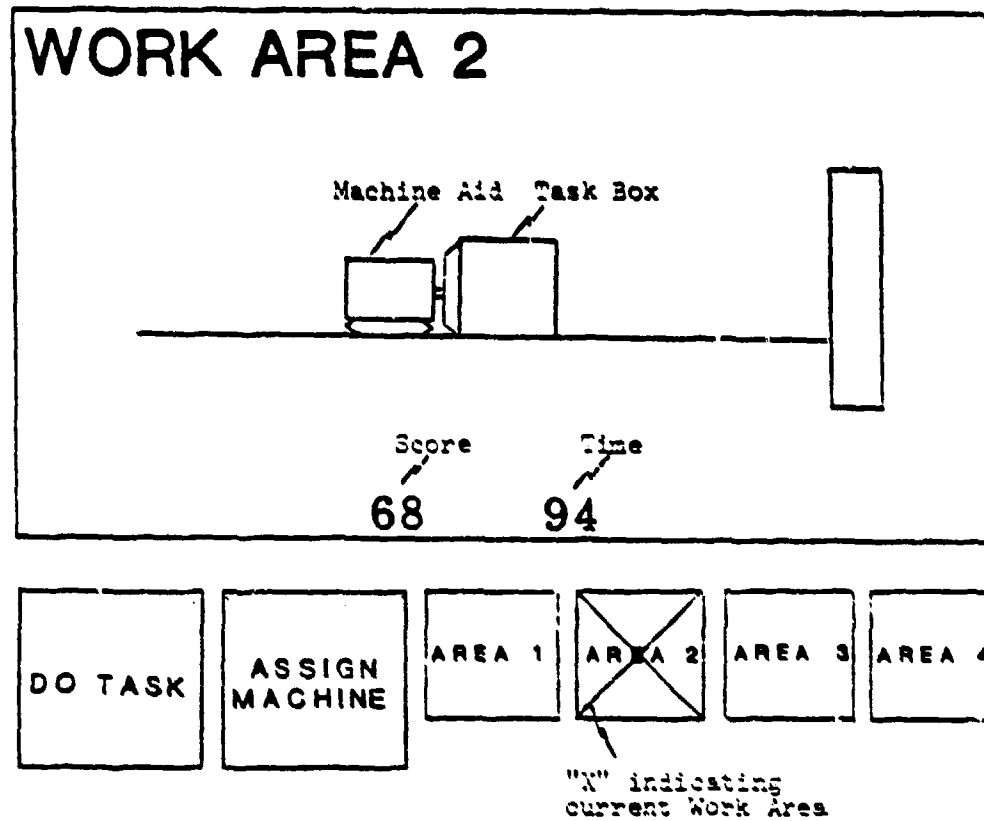


Figure 23 : Diagram of the SUPER playing field

From the probabilities at each point in the tree the computer can determine the optimal strategy. However, the results from a decision tree which looks three steps into the future will be optimal only if the experiment is limited to three steps. To include more and more possible events in the analysis, the decision tree must be extended more and more steps into the future. In the limit, as the number of steps becomes infinite, the decision tree will incorporate all possible future development. Unfortunately, an infinite decision tree will have an infinite number of branches and the computation of expected RPT becomes impossible.

As the number of steps gets large, the incremental benefit of looking one more step into the future will decrease. It is possible to find a number of steps N that is sufficiently small to allow for computation of the "best" path, but which is sufficiently large to approximate the infinite tree. These approximately best paths or strategies are referred to as the "N-step optimal".

This model was validated and an appropriate number of steps was determined by running simulations. In cases where a low N model was found with the same effectiveness as a high N model the low one was chosen because of decreased computational time. In general, a six-step model was found to be most effective for determining the "best" strategies in this study.

5. Measuring Human Performance on SUPER

Human performance was measured by isolating each decision made by the subject and recording the game conditions at the time of the decision. The best-choice decision based on the optimal N -step strategy was then computed for the given conditions. The fraction of subjects' decisions which agreed with the computed best strategies was the resultant performance measure in this case. By analyzing each decision individually, separate performance measures can be computed for the different types of decisions. By comparing human and optimal performance measures it is possible to isolate the causes of human sub-optimality.

As suggested earlier, by employing a single dimension reward we expect to remove much of the nonlinearity and suboptimality in operator inferences. Another cause of sub-optimality, however, may be the operator's internal misrepresentation of the task parameters by the subject. While all task parameters were presented to the subject, in general he will accurately retain only some fraction of the information given. Some information he will forget and some he will make up to take the place of forgotten information. Unfortunately the link between the information presented and the information used by the subject to make decisions is not apparent.

Methods of presenting task parameters were varied in hopes that one would prove especially effective. Some improvement in retention was seen from relating parameters to stories and allowing subjects time to study them. Each subject was debriefed after the experiment to check his recall of parameters, strategy, and evaluation of performance relative to his strategy. Errors in recall were used as evidence of internal misrepresentation.

To try to further limit internal causes of sub-optimal behavior a task familiarization and training period was given each subject prior to the experiment. During this period the subject became accustomed to the computer display and his score was displayed as feedback for this practice time. The subject was exposed to only one task class at a time in this task familiarization period, thereby removing from consideration the decision to change work areas. In this way the subject was forced to concentrate on the

relative merits of performing the current task by machine versus doing it manually. He was also forced to sit through the machine service time and time between task arrivals, items he might otherwise ignore while in the SUPER experimental situation. This practice served to bring subjects closer to steady-state learning. Once the experimental period was begun the score display was removed so that the subject could not use changes in the score (such as decreases due to holding costs from a task appearing in another work area) as a strategic indicator.

Subjects in SUPER experiments had three demands on their mental resources: they had to receive information from the SUPER display, they had to make decisions, and they had to implement these decisions. If the "rate" of an experiment, measured in terms of task arrivals and task completions per second, was increased it would seem that subject performance suffered because less time was available for decision making. A small series of experiments was conducted in order to determine a "fair" or "comfortable" experimental rate where subjects were not rushed in making decisions.

Three subjects were brought into the laboratory for three sessions each. In each session three experiments were conducted. These experiments were varied in the number of work areas used (called the scenario) and the experimental rate. Three scenarios employing two, four and eight machine aids were used, as well as three rates, fast, normal and slow.

It was expected that the "fast" rate would cause a decline in performance, but it should be noted that the "slow" time caused a decline as well. Subjects attributed this to boredom. They felt the excess time, and reported that long periods of inactivity made them forget what work areas they had viewed recently, thereby hindering their search strategy.

There were absolute differences in performance from scenario to scenario attributable to the differences in task environments. The two-work-area scenario showed better performance than that with four areas, because the search strategy required was less complex. The eight area scenario also showed good results because all tasks within that scenario were identical, so score depended on search strategy rather than on rewards and cost.

6. Experimental Results

Operators in SUPER faced two basic decisions - when to change work areas and whether to do tasks manually or assign machines. Our interest was chiefly concerned with the latter. Approximately seventy percent of all decisions faced by subjects, however, involved the question of whether to change work areas. A few general results regarding search strategy are therefore presented.

Transition Time. The importance of the decision to change work areas was dependent on the transition time. If the time was short, the effect of sub-optimal strategy was small because errors could be corrected rapidly. Operator decisions reflected this fact. When transition time was very short operators reported in debriefing that they only employed very simple search algorithms. When the time was lengthened, the subjects reported much more complex strategies.

Human Processing Limitations. A subject's search strategy depended less on the current condition of the experimental system than it did on the search pattern he chose to use. These search patterns or algorithms seldom took into account more than the last areas visited and where work was most recently found. Many simply pushed the work area controls in sequence,

demonstrating how crucial the role of physical design may be. The exception to this was for work areas with very low arrival rates. Once these were discovered, subjects tended to skip these areas in their search patterns.

Machines That Are More Productive Than Their Supervisors. In assigning machines which could perform tasks faster than their supervisors subject performance was quite close to optimal with deviations attributable to search strategy. When the subject was given a machine that was highly productive at one task class but not at others, he assigned the machine there and performed the others himself. If the machine was efficient at all task classes, the subject would use it for virtually all the work.

Men and Machines with Comparable Abilities. Situations where men and machines are interchangeable are not uncommon. An airplane cockpit, as mentioned at the outset of this report is one example. When subjects were given machines of ability similar to their own they performed quite well. A standard strategy was to search for a task and then assign a machine aid to it. If no machines were available the operator began the task himself until a machine became available. The N-step model generated an almost identical strategy.

Low Productivity Machines. In low productivity experiments subjects were given two machine aids, one with the same productivity as the operator, the second with productivity scaled down by a factor of X, with X ranging from 1 to 10. Subjects were told they did not have to use the less productive machine if they didn't want to. Yet, as demonstrated in Figure 24, subjects used the less productive machine far more often than the optimal model did.

When questioned, subjects said they felt they would be wasting some of their resources if they didn't use the machine. A common evaluation of behavior was that if a low productivity machine was assigned in a work area not often visited, temporarily at least, the number of areas to be searched was reduced. Subjects also felt it was better to have as many tasks as possible being worked on even if that meant using an inefficient machine.

The Effect of Machine Wages on Strategy. Increasing machine wages should make a machine less desirable to use. Decreasing productivity should have a greater inhibiting effect because it not only spreads task rewards over a longer time but prevents new tasks and their potential rewards from appearing in the work area. This was not, however, borne out by the experimental results. Surprisingly, disinclination to use machines as machine wages increased was much more rational than in the case of decreasing productivity. However, as can be seen in Figure 25, subjects were not able to be quite as discriminating as the optimal model.

It is unclear why machine wages should have a greater impact on human behavior than productivity. It may be because in the task familiarization period the operator could see how fast his score started to drop when he used an expensive machine, i.e., negative feedback was more apparent and immediate. It may also be that subjects found cost a more tangible quantity than productivity. Unfortunately, machine wages are not one of the easily adjustable quantities in system design. Also, when an operator is an employee he is not the person who must pay the machine wage. Therefore machine wage may have only a very limited effect in the real world. When operators have to take on responsibility for the cost of their machines they may become much more inhibited than otherwise.

7. Conclusions from SUPER

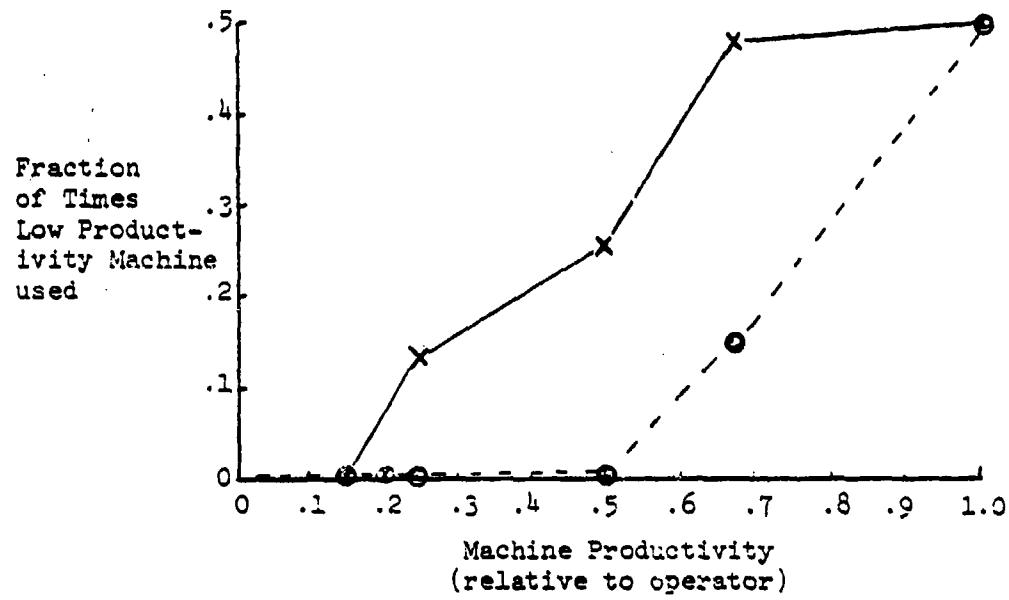


Figure 24 : Fraction of all decisions to deal with a task where the low productivity machine was used - for decisions made by
 —x— experimental subjects
 --●-- the optimal model

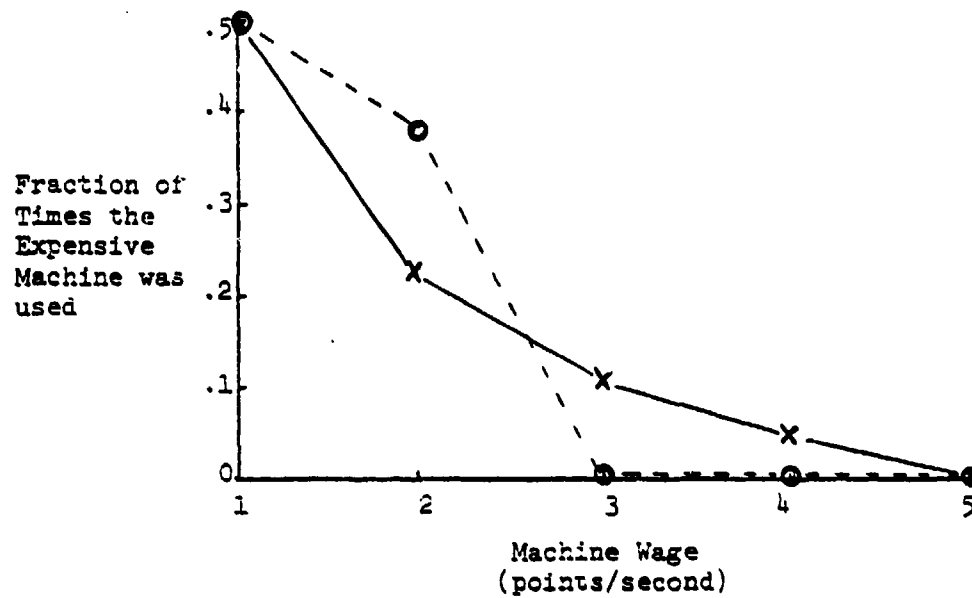


Figure 25 : Fraction of all decisions to deal with a task where the expensive machine was used -- for decisions made by
 —x— experimental subjects
 --○-- the optimal model

A system designer concerned with the efficiency of an operator's use of machines should not be worried that humans will usurp jobs that should be performed mechanically. In fact, this study indicated that the opposite is likely; the operator will use the machine more than is optimal. Making operators aware of the true costs involved in machine use (maintenance, depreciation, capital investment, etc.) may decrease this tendency.

The experiments in this study dealt specifically with machine aids, and it is not clear whether the results are generalizable to the assignment of any tools, or to workers with variable productivities, or to tasks with changing requirements. Certainly machine productivity could be more deeply explored, especially in regard to the probability of machines successfully completing tasks. In this study machines did not fail; once a task was assigned to a machine it was always completed. In real life situations this is not always the case, and the effect of possible machine failure on operator performance is certainly an area for exploration.

Further, in this experiment the only requirement to put a machine to work was the push of a button. Real machines usually require set-up time, and this could be included in future studies.

Finally, our experiments used finite task queues. A new task could not appear in an already occupied work area. An implicit cost of leaving a task unattended in this study is that the rewards of possible new tasks in that area are foregone. By employing infinite queues where tasks "line up and wait" the costs of using a slow machine might be reduced and subject performance improved.

III. ILLUSTRATIVE EXPERIMENTS ON SENSING/DISPLAY AIDS TO THE SUPERVISOR

* This section draws on Ranadive and Sheridan [1981], Ranadive [1979] and Deghuse [1980]

The experiments described below are concerned with the "affecter" loop in Figure 2 (loops 1 and 2) and the "automatic" simulation/cognition aid, loop 8, which supplements the feedback from the remote process itself. Section III-A deals with the problem of constrained bandwidth in loops 1 and 2, one of the key reasons for having supervisory control of remote manipulators and vehicles. Section III-B deals explicitly with loop 8 from the viewpoint of supervisor simulation aiding. Section III-C tells of a special way to extend and enhance sensing and display by letting a low-level computer build up a visual image, thereby aiding the supervisor's memory (which method is curiously dependent on motor control, a kind of inverse of section II-C. Section III-D gives an example of a well known display enhancement technique called a predictor display - one which promises to be very useful in supervisory control. Section III-E and III-F deal with a special class of sensing/display aiding, but one which will assume greater and greater importance as automation increases, namely, that of aiding the supervisor in detecting and locating failures.

A. Bandwidth Limitation In Telemanipulation: The Framerate, Resolution, Grayscale Tradeoff

1. Introduction

One reason for using supervisory control in space, in the deep ocean or indeed on terra firma is because the communication between human operator and a remote system is severely constrained, that is, the bandwidth is limited. For teleoperation in deep space one good reason is that radio energy is dissipated over the long distance to be spanned. If an electrical cable is used for deep ocean operations there need be no such problem, but there are problems of the tether becoming a large drag on the submersible vehicle and/or getting tangled up in structures that one wishes to inspect remotely. To avoid the latter problems one may employ acoustic communication. Even if a tether is dropped from a surface vessel down to within a few hundred feet of the submersible, acoustic transmission for the remainder of the distance can circumvent the problems cited above. However, because of sound energy dissipation this can only be done at the cost of having to reduce the bandwidth considerably relative to that for a wire (tether).

On terra firma as well as in space or undersea another bandwidth reducing factor may be that the operator may have to time-share his attention.

Thus one is left asking, for a given fixed communication bandwidth, how best to trade between the three variables of frame-rate (frames per second), resolution (pixels per frame) and grayscale (bits per pixel), the product of which is bandwidth (bits per second).

2. Experiments

These tradeoffs were studied by Ranadive [1979] in the context of master-slave manipulation. Experimental subjects were asked to perform two remote manipulation tasks using a video display as their only feedback while using our Argonne E-2 seven degree-of-freedom servo manipulator (in this case with force reflection turned off). Figure 26 illustrates the experimental situation.

The first task was to locate a nut on a fixed bolt or knob and take it off by unscrewing it. (We abbreviate this task "TON" for take-off-nut). The second task was to pick up a cylinder and place it sequentially within the bounds of three fixed squares on the table which were numbered 1,2, and 3, where the order of the placement, e.g. 3-1-2, was randomly drawn for each new trial. (We abbreviate this task "1-2-3"). Performance on each task was simply defined as the inverse of the time required to do that task correctly, and combined performance was the average of these inverse times.

The video display was systematically degraded with a special electronic device which allowed frame-rate to be adjusted to 28, 16, 8 or 4 frames per second, resolution to be adjusted to 128, 64, 32 or 16 pixels linear resolution and grayscale to be adjusted to 4,3,2 or 1 bits per pixel (i.e., 16, 8, 4 or 2 levels of CRT intensity). Figure 27 shows the effect of resolution reduction.

Two subjects were used, both engineering students. They were trained for 10 hours in all combinations of display tasks and visual variables. When subjects first saw the video pictures with which they had to perform remote manipulation tasks, they refused to believe that they could succeed. Much to their surprise, however, they discovered that they were able to perform with a considerably degraded picture. During the data collection phase of the experiment subjects were allowed to practice on each display combination until "ready".

The data collection runs were ordered so that two of the three video variables were kept constant while the third was varied randomly among the levels for that variable. Ten times were collected (ten trials were run) for each combination (each data point).

Figure 28 shows the results. On the top row are shown the performance effects of frame-rate, resolution and grayscale while holding the other variables constant. Note that for frame-rate beyond 16 frames per second improvement depends on resolution and grayscale; performance improves smoothly for increases in resolution; for grayscale there is no improvement beyond 2 bits if the frame-rate is high enough.

On the bottom row constant level-of-performance tradeoffs (in this case using the TON task only) are shown for each of the three pairs of video variables. These iso-performance curves (solid lines) are compared to iso-transmission lines, i.e., combinations of the two parameters which produce constant bits per second. It is seen that there is a remarkable correspondence. This means that for this experiment, and within the range of video variables employed, man-machine performance corresponds roughly to bits per second of the display, regardless of the particular combination of frame rate, resolution or grayscale.

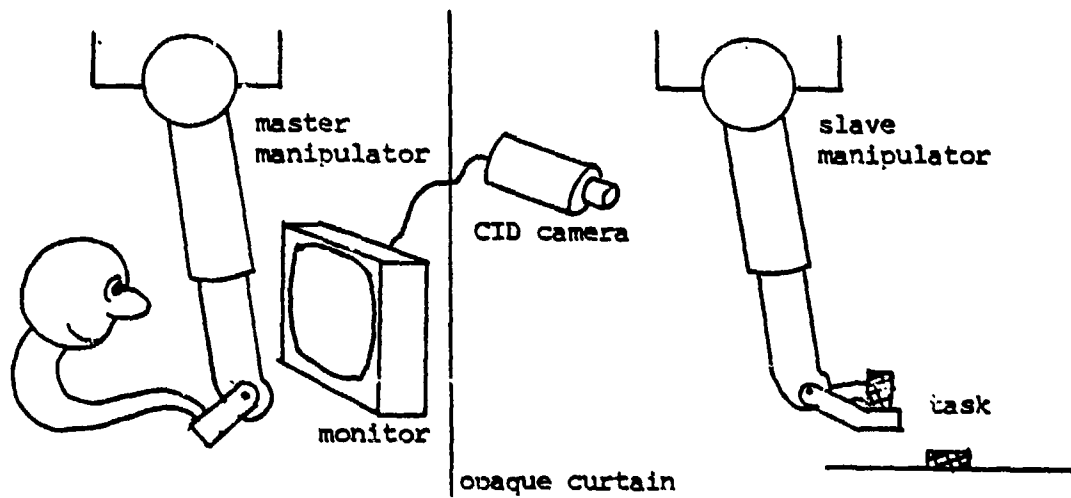


Figure 26 . Experimental configuration for Ranadive frame-rate, resolution, grayscale tradeoff experiments



Figure 27 . The same picture at various degrees of resolution (pixels)

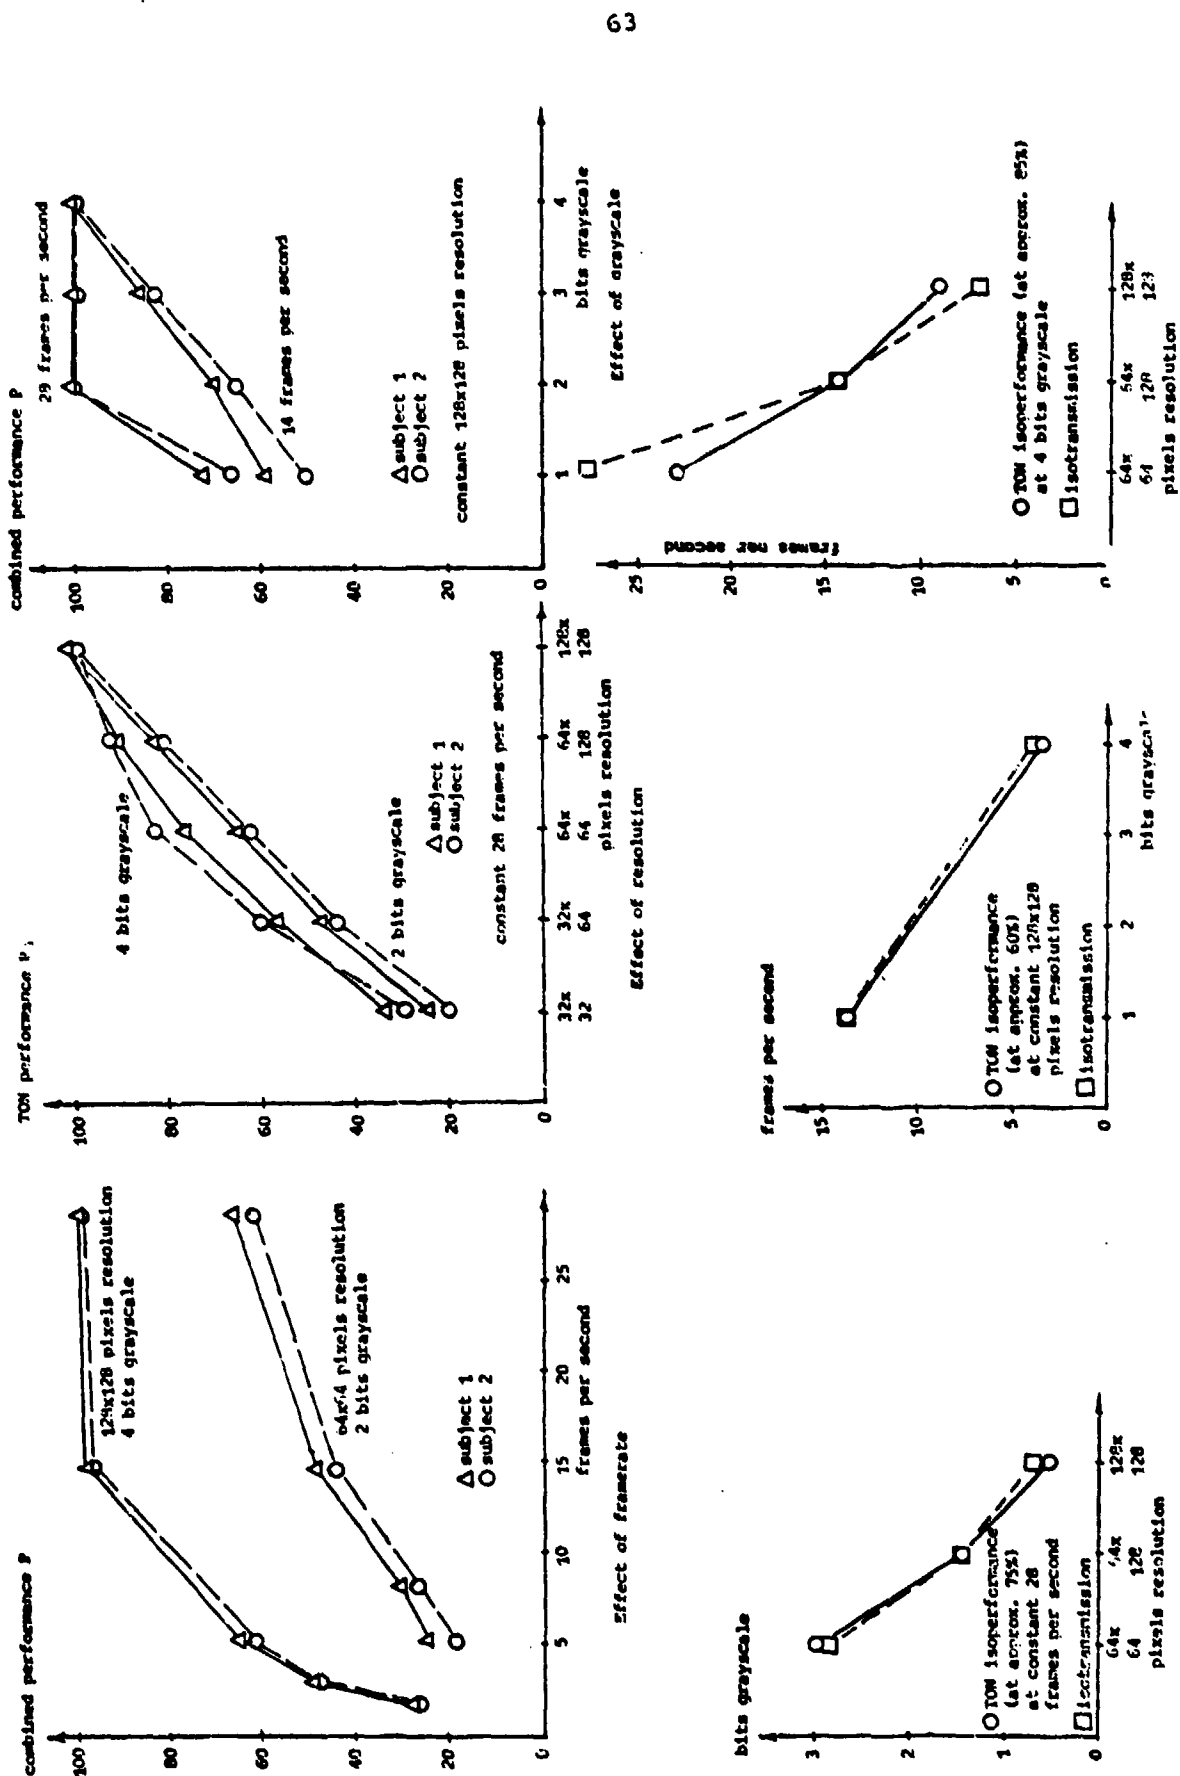


Figure 28. Results from Ranadive experiment. The top row shows as one of the three variables is adjusted, the other two are fixed. The bottom row shows isoperformance and isotransmission tradeoffs for two variables at a time, the third held constant.

Another result, though not tested systematically, was that subjectively much more noise appeared on each video picture at the slowest frame-rates than at faster frame-rates. It is believed that this was due to visual-psychological smoothing rather than anything electronic occurring at higher frame-rates.

Assume limited-bandwidth acoustic transmission is to be used as the means for communication between a human operator and an undersea teleoperator. It is clear that video will pose a far greater bandwidth requirement than other signals, probably more than all other signals combined. Therefore it seems reasonable to allocate fixed channels to other variables as required, then to reserve the bandwidth remainder to the combination of video signal requirements. That is, framerate, resolution and grayscale would not each have fixed bandwidth allocations; rather, provision would be made to trade-off between these as required, retaining their product as close as possible to the maximum.

To make this idea more understandable, assume that a given human operator of a teleoperator needs to get an accurate picture of a static object. He would like high resolution and sufficient grayscale, but frame-rate could be anything. In contrast, suppose the operator needed to monitor the speed at which a well known object moved against a background. Only enough resolution and grayscale would be necessary to get a good definition of what is object and what is background, but frame rate would have to be high. Either condition could be obtained by adjustment.

3. Operator-Adjustable F-R-G Tradeoff

Deghuse [1980] used an experimental computer-based aiding device which allowed the operator to make this three-way adjustment in-situ, i.e., he could modify the F-R-G (frame-rate vs. resolution vs. grayscale) tradeoff himself while performing a master-slave manipulation task of the type performed in Randive's experiments. In particular, the same master-slave manipulator was used with the force feedback turned off, and the same take-off-nut (TON) task, but a many-peg-removal task was chosen instead of the 1-2-3 task. Scoring was inverse time, the same as in Randive's experiments.

Four subjects were trained about ten hours each on each of the tasks. The same levels of the video variables were used as before. Two maximum bit rates (products of frames per sec, pixels per frame and bits per pixel) were used, one being 11500 bits per second, the other twice that or 23000 bits per second. The subject used three keys to call for any F-R-G combination, up to the maximum. When one factor was increased the other two were automatically decreased to keep the product at the constant maximum. Each subject, for every combination of task and maximum bit rate, performed both with and without the in-situ tradeoff adjustment capability. There were four trials for every cell of the experimental design. Data were analyzed by analysis of variance.

As might be expected, the use of the tradeoff control was significant ($p < .05$). Further, both the task main-effect and the task-subject interaction were significant ($p < .01$), a result not particularly surprising. What was more surprising was that the two maximum bit rates did not produce significantly different performance.

There was much variability in performance due simply to the fact that the visual interpretation time was extensive, and the real-time continual decision task of how to set the F-R-G combination added to this. It is believed that the means of making this adjustment could have been better "human engineered", and that this would have reduced variability and improve performance. Similarly the lighting was seen to be a critical factor, where amount of light affected grayscale adjustment and shadows provided important cues.

A principal result of this study was confirmation that with some training and some patience an operator can remove a nut with a remote manipulator using video of only 10 bits per second and with no force or tactile feedback. From the results an important special use of the adjustment became apparent in this case, namely to periodically but briefly increase resolution and grayscale at minimum frame rate in order to get confirmation that the peg was in the hole, or that another critical task phase had been achieved.

Use of this device is an important aspect of supervisory control, where the computer aid mediates the operator's instructions to provide, in this case, the best display (rather than control per se). This is loop 8 (of Figure 2) working in conjunction with loop 2.

B. Computer Graphic Manipulator Simulation For Planning, Training And Real-Time Feedback

* This section is based on Winey and Sheridan [1983] and Winey [1981]

1. Introduction and Objectives

The objective of these experiments was to explore the development and use of a flexible computer-graphic simulation of a master-slave manipulator, a simulation which could be operated by a person in real-time and provide both visual and force feedback. Such a simulation would allow the operator, in effect, to see and feel his machine and work environment from any viewing or probing angle and at any scale, and to try out a variety of tasks, designs and control strategies without commitment to real hardware and associated costs and risks.

A realistic computer-graphic simulation can be of value both for experimentation and for operator training. Environments such as deep ocean work areas, which cannot easily be recreated physically in a laboratory, can be simulated on a computer instead. Further, by using the computer to control the simulated manipulator, it is possible to vary the kinematic and dynamic properties of the manipulator so that it can be used to simulate many types of manipulators. Several specific applications of the manipulator simulation are listed below:

2. Testing of Control Systems

Building a control system for a remotely supervised vehicle can be both expensive and time consuming. With a new controller, there is always the risk of instability and failure which can result in damage to the hardware. When a control system is being developed for a one-of-a-kind prototype vehicle, the money and time lost in a failure can be disastrous. A computer simulated prototype controller can be changed quickly and easily, and failure generally involves little or no risk. A simulation can be run in and around instability in order to collect failure data without damage to hardware.

Supplementing of Operator Visual Feedback - A simulation may be especially helpful in remote work environments from which it is difficult to obtain pictures quickly or accurately. For example, as mentioned earlier, when unmanned vehicles are used in the deep ocean to avoid tethering problems they may be connected to the control system via an acoustic link. Such a link is subject to very limited (low bit rate) transmission; a television picture relayed this way must have either low resolution or low frame rate or both. The simulation, however, can be generated and updated in real time with only a small amount of data. All that must be known continuously is each angle of the manipulator's seven degrees-of-freedom. By superimposing a rapidly updated simulation of the manipulator on a slowly updated, but high resolution, television picture, the human operator is provided both the movement and resolution information that he needs.

Rehearsal - When an operator is required to perform a dangerous or delicate task in which a mistake could be harmful to himself, equipment or the task itself, it may be useful for him to practice the task on a simulator

until the operator feels confident of his performance. If the manipulator is computer-controlled, the computer can monitor each of the practice runs and, when a satisfactory run has been achieved, the computer can be told to duplicate all or some part of it.

Humanizing Man-Computer Interaction - In addition to aiding the display of information from the work area to the operator, the computer graphic simulation can be used to accept information from the operator and to process it. The master manipulator can be used as a means of giving commands to the computer with the simulator serving to acknowledge understanding of the command.

The manipulator-plus-graphic-display can be used as a three-dimensional digitizer. Being capable of force-feedback it could assist the operator in locating reference points for such input, where some information, e.g., some attribute of position, is fed back to the operator as a force signal. The preceding are just a few examples of the practical applications of a computer graphic simulation of a remote manipulator.

3. Developing the Simulation

The major pieces of computing equipment used in the simulation were a PDP 11/34 computer running the RSX-11M operating system, and a Megatek 7000 Vector Display Processor with 3-D hardware rotate and 4096 lines resolution. The manipulator used as a master and simulated as a slave was the Argonne E-2 master-slave manipulator. It has seven degrees-of-freedom of motion and full bi-lateral force-reflection. Its electronical coupling allows it to be interfaced to the computer through an AN5400 A/D converter. Figure 29 shows the slave portion of the manipulator which was simulated.

These were stored in the computer using standard point-connectivity data. The manipulator arm was described in three separate pieces - the shoulder, forearm, and tongs - each treated as a separate entity. Each of the geometric elements was stored in an unrotated reference frame, and given a corresponding rotation matrix containing the transformations required to move that element from the reference frame to the desired location. Environmental objects for manipulation were further defined by "touching conditions" which described whether or not the object was in the manipulator tong's grasp.

Two types of objects and touching conditions were tested in the simulation experiments. The first object used was a sphere with a program specifying spherical touching conditions. When this proved successful, a stepped rectangular peg was also tested. Two sets of rectangular touching conditions were defined, one for the main body of the peg and one for the stem.

Various dynamic and static properties for the objects, such as gravity, viscous drag, elasticity, and conservation of momentum, were included in the simulation and could be modified. This allowed great flexibility in simulating the environment within the computer. The manipulator and objects could be enclosed in a rectangular room. When a moving object collided with a wall of the room, it rebounded. This served to keep the objects within reach of the manipulator, as well as demonstrating conservation of momentum.

The simulation could be displayed from any viewpoint, and the viewpoint could be either stationary or in motion. The display could also be scrolled and zoomed so that any portion of the display could be observed in detail.

A simple submarine vehicle simulation was added, and made to be capable of the same functions as the manipulator simulation. The manipulator was mounted on a vehicle which could be controlled in six degrees-of-freedom. The position of the vehicle was entered through an installed common block, which allowed the vehicle to be controlled by a secondary program.

Force Feedback - Two applications of force-feedback were introduced in the simulation. When a simulated object was gripped by the simulated manipulator, force-feedback was sent to the actual master. This was sufficient to keep the simulated slave tongs open to the width of the object. The resulting sensation felt by the operator was that of an actual object within the tongs. The second application used force-feedback on all joints of the manipulator: A three-dimensional elastic surface was defined (Figure 30). The neutral surface (no force applied) was assumed to be relatively flat, eliminating the need to calculate the surface normal. Instead, the surface normal was assumed to be vertical.

Different elastic coefficients were assigned to various locations on the surface. Force-feedback to the master was generated proportional to the penetration distance beyond the neutral surface and the stiffness, so that the surface could be felt when it was touched. To provide a visual indication a gridwork approximation of the surface was displayed on the graphics terminal. To aid the operator in perceiving depth, the contour directly under the manipulator was displayed in darker linework. As the manipulator penetrated the surface, the contour deflected. If the surface was soft, deflection only occurred in the neighborhood of the penetration. If the surface was stiff, a larger portion of the surface deflected.

Depth Indicators - The major difficulty in using a video terminal for a manipulator simulation display is the lack of depth perception available to the viewer. When the simulation was first developed depth information was transmitted using traditional orthographic projections. This approach caused a few coordination problems and operators could become confused as to which view was front and which was side. As a result two other types of depth cues were tested to try to improve the operator's depth perception.

Experienced manipulator operators often rely heavily on shadows for depth information (Figure 29), so shadows with the source of illumination directly overhead were chosen as a second depth cue. The shadow was cast on an imaginary horizontal floor 50 inches below the manipulator's shoulder. Walls displayed on the screen were helpful in correctly orienting the shadow.

Unfortunately, to be understandable both the orthographic and shadow cues required extensive prior knowledge of the environment. The third depth cue display installed in the simulation, which did not require such environmental knowledge, was a proximity indicator which showed the absolute distance between the tongs and the object. This was displayed by a simple line on the terminal screen. The display was designed so that the length of the line was on the same scale as the display and was present any time the

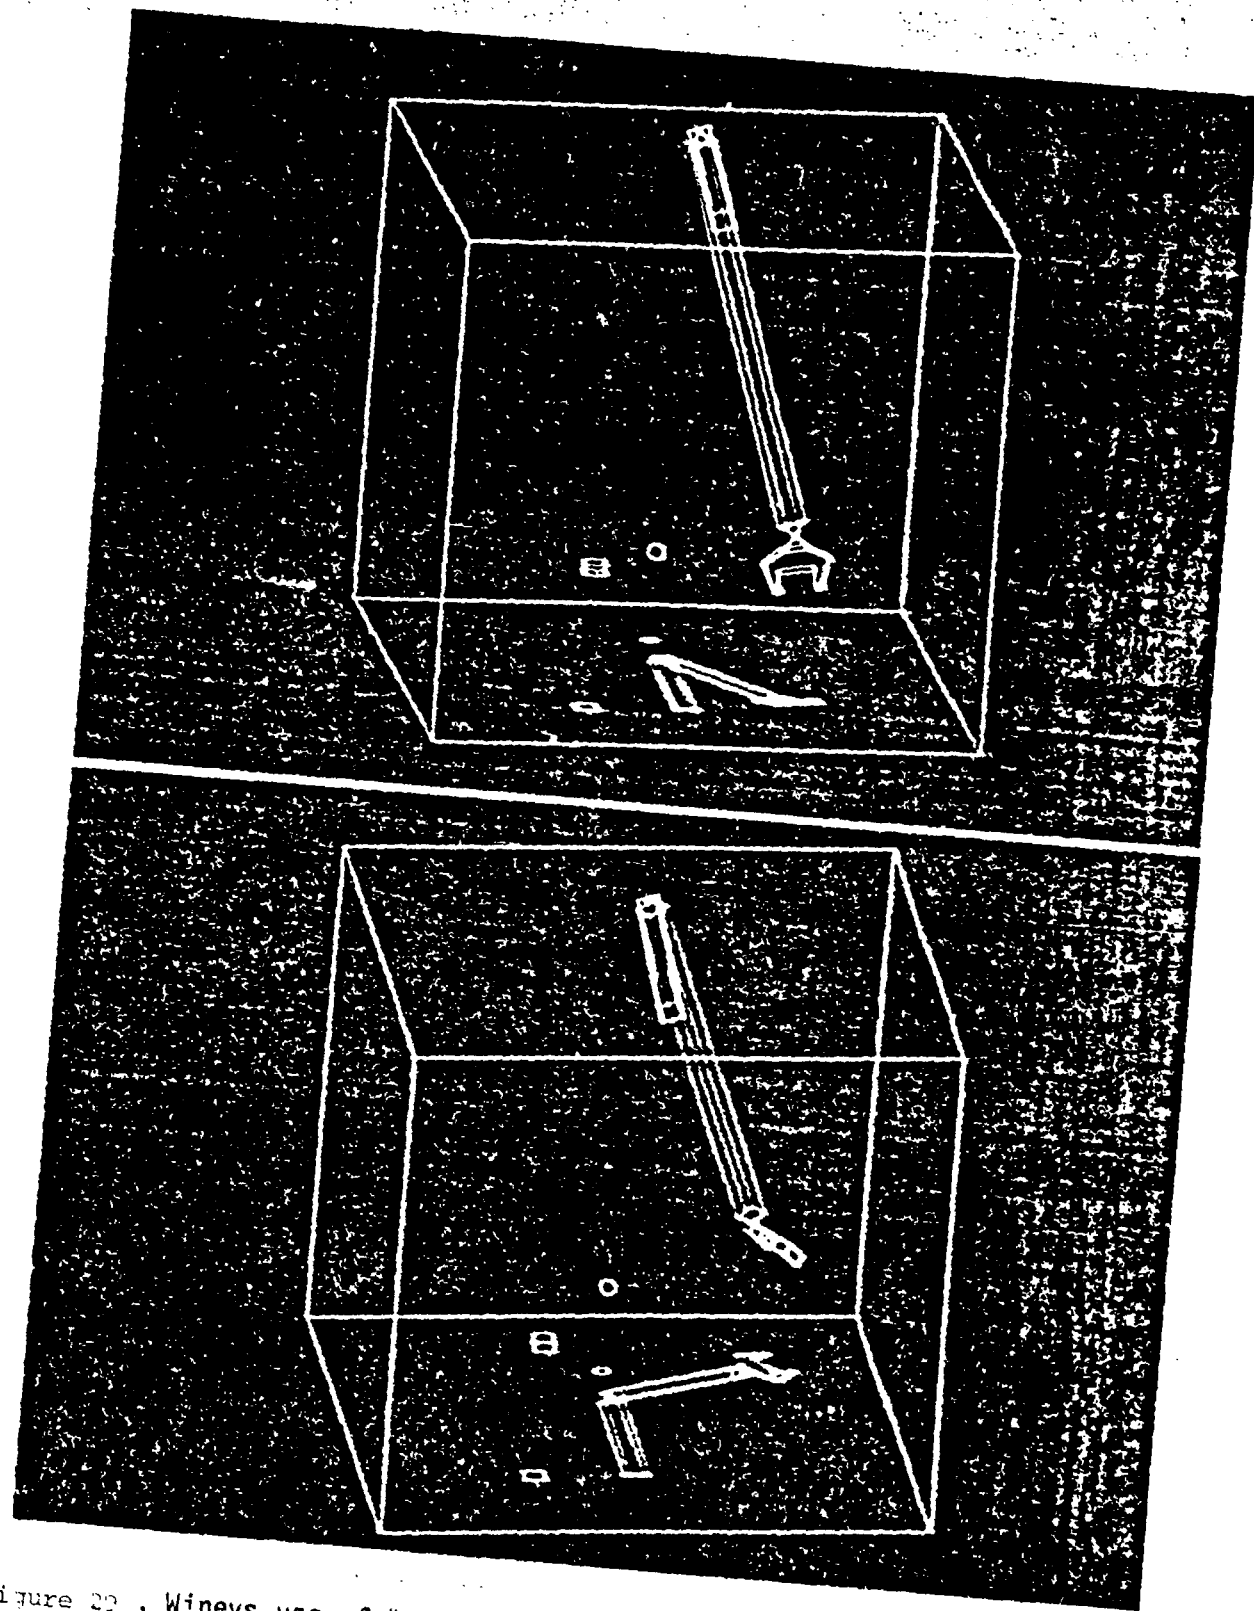


Figure 22 . Wineys use of "shadows" and surrounding box to teleproprio-
ception

object was within 24 inches of the tongs. The indicator was ten times less sensitive at longer ranges. In real-life situations a proximity detector and indicator could be implemented using a sonar device on the tongs of the slave manipulator.

4. Experimental Measurement of Operator Performance Using the Depth Displays

Three depth indicators, front plus side (orthographic projection), shadows, and absolute proximity in addition to the control were tested on five subjects. Two types of tasks were designed. The first required the operator to reach out and grasp a simulated two-inch, stationary sphere. The time it took the subject to grasp the sphere after the display was flashed on the screen was recorded. In the second task the conditions were the same except that the sphere was in motion. The sphere followed an orbit path about the surface of an ellipsoid. It was hoped that the moving task would highlight any coordination problems associated with the depth indicators.

The subjects were tested on a series of display types with both moving and stationary objects. The positions of the objects were selected so that they were always within reach of the manipulator and did not coincide with any real object which would obstruct the manipulator. The display types (depth indicators) were intermixed because subjects tended to become bored with repetitions of the same display. The order of the display was, however, kept constant, so subjects knew which type of display would be next. The positions of the sphere were arranged such that the average distance between successive positions was the same for each display type.

Four subjects each performed 80 to 90 repetitions of each combination of the four display types and two tasks. Two of the subjects had prior experience with the E-2 manipulator and the graphic simulation. The third had experience with only the manipulator, and the fourth had no prior experience. All subjects learned to use the displays in one to two hours. Learning curves were recorded for each subject to insure that performances plateaued before regular test trials began.

The analysis of how well each depth indicator performed is summarized in Figure 30. Clearly the proximity indicator with no additional display (the experimental control) gave poor results. So that it would not overwhelm the other three display types, the proximity indicator alone was left out of the statistical analysis.

A three-way analysis-of-variance was then performed. The results showed marginally significant differences between subjects, display types, and tasks. On both stationary and moving tasks, the front and side orthographic projections showed the best performance. Three of the four subjects said they preferred this display because it gave the clearest detail. All said it presented only slight coordination problems.

The use of shadows yielded the second best response times. All the subjects felt the shadow gave them the best perception of the object's position in the environment. The main difficulty with the shadow depth cue was that the manipulator's shadow tended to obscure the object's shadow when the two were in close proximity. Further, and although the result was not

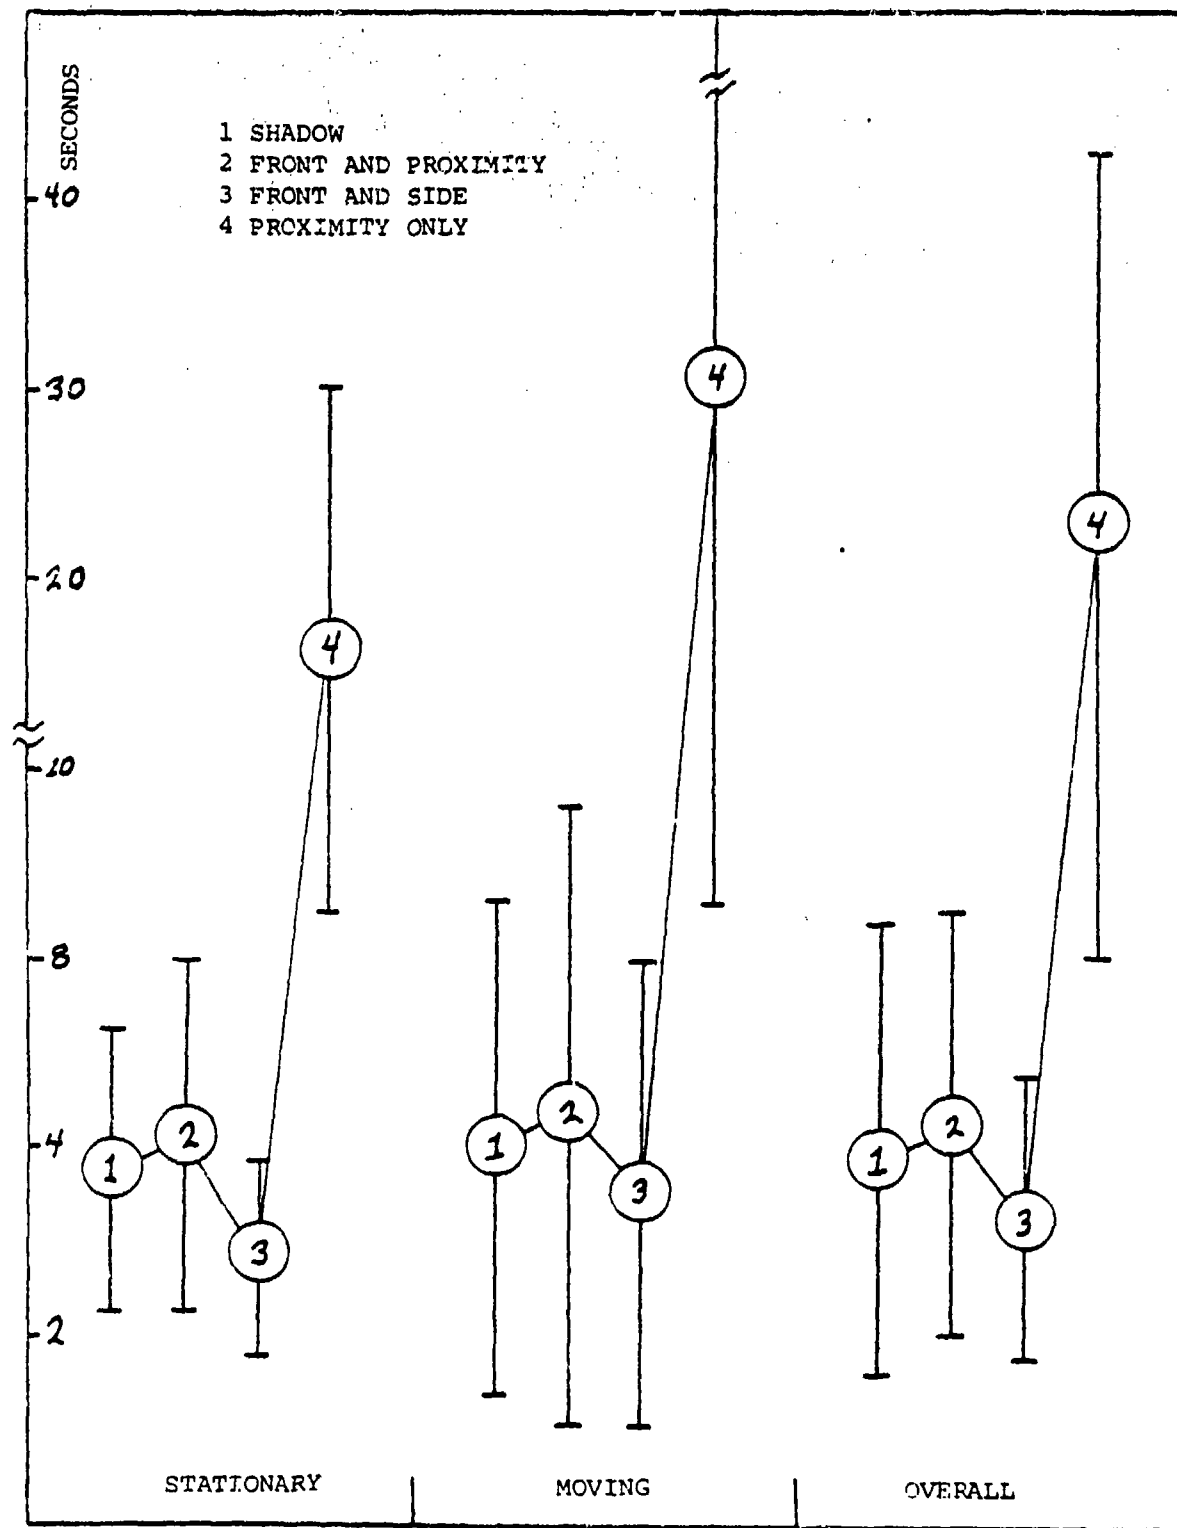


Figure 30 Evaluation of Depth Indicators

statistically significant, the shadow indicator seemed less affected by the motion of the object than by the front and side views. With larger or faster motions this could prove a significant advantage.

One subject preferred the front view of the manipulator with the proximity indicator because of its simplicity. His response times using this cue were close to those of the shadow. In order to obtain depth information from the proximity indicator the operator had to move the manipulator and watch the indicator's response. This caused the subjects little trouble, although occasionally they would search in the wrong direction initially. The proximity indicator is easily implemented in a simulation, but because it provides quite limited depth information it is probably not suitable for a complicated task.

Although there were differences in the effectiveness of the three depth indicators, these differences were small compared to the corresponding times required to perform the task. It is difficult to say with certainty which of the three indicators was best; each had advantages and disadvantages. Based on the experimental results, we suggest that the front and side views be combined with the shadow. The shadow provides an overall perception of the environment, while the side and front views elaborate the detail. An alternative in practical applications would be to allow the operator to select the view with which he is most comfortable.

5. Evaluation of Simulated Force-Feedback

One feature of the simulation which appeared quite promising but was not systematically evaluated was the use of actual force-feedback from the simulated forces. Force-feedback was generated when an object was grasped, giving the simulation a strong feeling of reality (Figure 31). Evidently the feedback was used by subjects in the simulation in two ways.

The first use of force-feedback was confirmational; it let the operator know when an object was grasped by the manipulator. If the object happened to slip from the tongs, the loss of tactile feedback let the operator know immediately.

The second use of feedback was quantitative. The force-feedback could convey information about the object, such as stiffness or weight. Surfaces ranging from extreme rigidity to a consistency approximating foam rubber were simulated. The main difficulty with simulating a variety of surface types came in the cycle time of the simulation. The softer the surface, the deeper the manipulator tongs would penetrate under a given applied force. The larger penetration motion required the manipulator to move greater distances on each program cycle. As a result, the simulation program had difficulty updating the manipulator quickly enough. Despite this problem, it was possible to simulate a wide range of surfaces with good response. Further testing and refinement of force-feedback would seem a promising area for further research.

6. Conclusions from Simulation Experiments

The manipulator simulation is believed to be a highly flexible and valuable tool for use with remote systems. The combination with environment

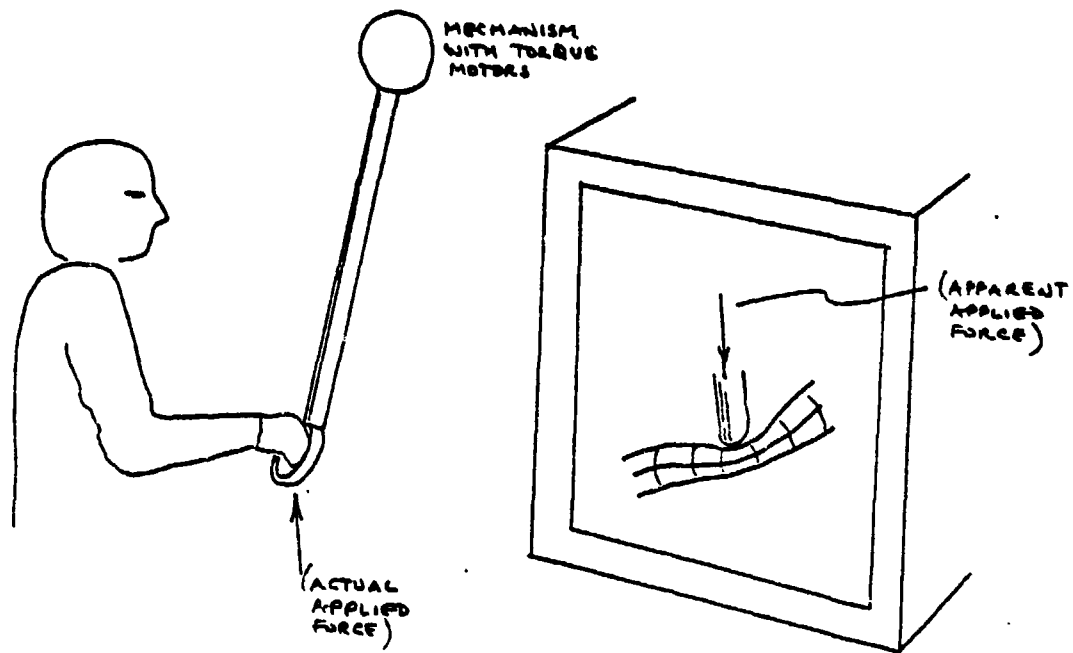


Figure 31. Winey's experiment with touch (force feedback from a computer model)

and vehicle simulation also shows promise, allowing the operator to know the arm-vehicle-object positions in real-time instead of requiring the interpretation of pages of computer output. Further, the visual feedback is compatible with both manual and computer control.

Having a real human operator with an actual master manipulator interact with a simulated slave manipulator and environment or vehicle proved very workable and suggestive of various uses. Future research might look to improving the reality of the simulated environment or to using the simulation to rapidly evaluate alternative control strategies for alternative (future) environments. This would be particularly useful in-situ, when the environment is uncertain and the cost of making an error in the soon-to-be-encountered real environment is high.

Study of collisions between the object and other portions of the manipulator appears promising. For example, this might allow the object to be pushed by the manipulator without being gripped. The use of force-feedback could be expanded, allowing, for example, the object to feel heavy when picked up, or simulating a reaction force when the manipulator is struck by a moving object. An additional area of application is in use and evaluation of touch sensors, i.e., arrays of force-displacement sensors on the skin of the gripper which allow differentiation of force in space as well as magnitude and time. A third is to evaluate the video overlay of a rapidly updated arm simulation with an actual slow frame-rate but high resolution video picture.

C. Blind Tactile Probing And The Inference Of A Computer-Graphic Picture Of Environmental Objects

* This section is based on the work of Fyler [1981]

1. Introduction

This section describes a novel means for tactile probing and discovery of the shape of an unknown object or environment. This technique offers promise for undersea operations where the water is so turbid that video is useless (and because high resolution acoustic imaging is as yet unavailable). It is the analog of a blind person probing in the dark by repeatedly touching at different points on an object or environmental surface in front of him and gradually building up a "mental image" of what is there, continually guiding his touching activity on the basis of what he discovers.

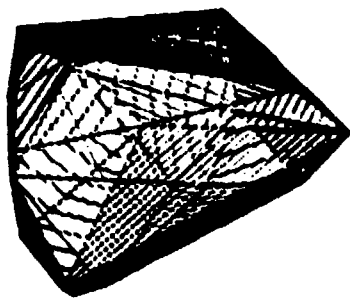
In performing "tele-touch" with a master-slave remote manipulator, if there were no dynamics and if force feedback were perfect it might be asserted that building up the necessary "mental image" would be no different than direct manual groping in a dark room. However every manipulator operator knows that is not reality; the master-slave manipulator itself is sufficiently cumbersome that one quickly loses track of where contact has recently been made and what the arm's trajectory has been. In performing tele-touch where a computer is determining the trajectory rather than a human operator's hand movements guiding a master, building up the "mental image" is still more difficult.

2. Touch Probe Display

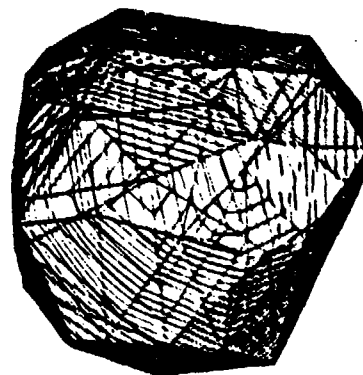
Fyler designed a unique touch-probe, a mechanical device which closes an electrical contact when it encounters a slight force from any direction. Then he programmed the 11/34 computer to determine and store the cartesian coordinates where any contact (touch) is made. He displayed on the Megatek screen, along with Winey's arm simulation, a projection from any viewpoint of cumulative touch points so stored. The operator can make no sense of such a display so long as the points are fixed. But the instant the image of points is rotated the shape and orientation of the one or more surfaces on which the contacts were established becomes immediately evident. What is a "mental image" in the case of direct manual grasping or touching becomes an explicit visual image. Figure 32 provides some (unfortunately static) examples of such displays.

As more points are added, the definition of the surface or object becomes more apparent. It helps somewhat to have the computer connect adjacent points with lines so that the best available "image" in three dimensions is a polyhedron and its planar projection is a polygon (or, of both front and back surfaces of an object are touched, two overlapping polygons). When rotation is effected the polyhedron immediately becomes evident. Rotation may be at a constant rate - usually around an axis near to or transecting the surface or object of interest - or may be controlled manually by a track-ball.

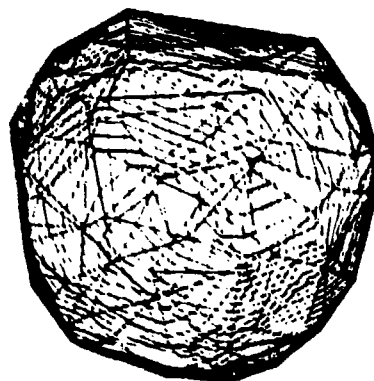
As contacts are made, points are added to the display, and what started out to be a polyhedron with few vertices and faces becomes a smooth



20 Points



50 Points



100 Points

Figure 32 . Random touch points on a sphere which generate a polyhedron.
As sphere is rotated the shape is easily perceived

surface, or a recognizable object. The first few contacts between the manipulator probe and environment are made more or less at random. However, as the polyhedron takes on form, it is evident to the operator where to place the next few probes to provide the most discrimination and not waste effort and time by probing in the wrong places.

Another display trick Fyler demonstrated was to put the polyhedron into the (Lexidata) raster-graphic display generator's look-up table in such a way that the orientation of any facet of the polyhedron is determined. Then by use of the look-up table he "illuminated" different facets of the polyhedron on the raster display as a function of the orientation of each facet - as if the sun or light source were at one angle shining on a polyhedron (Figure 32). Again the operator was provided a trackball, in this case to let him move the apparent light source to any radial position surrounding the object, the polyhedron in this case being fixed in orientation, not rotating.

D. Computer-Graphic Predictor Displays For Remote Vehicle Control With Transmission Delay And Slow Frame Rate

* This work was first described in Sheridan and Verplank [1978]

1. Introduction

Another form of computer-based display aid is the predictor display. This is a technique in which a computer model of the controlled vehicle or process is repetitively set to the present state of the actual system, including the present control input, then allowed to run in fast-time, say 100 times real-time, for some few seconds before it is updated with new initial conditions. During each fast-time "run", its response is traced out in a display as a prediction of what will happen over the next time interval (say several minutes) "if I keep doing what I'm doing now". The general technique is about thirty years old, has been much discussed in the human factors literature [Kelley, 1968], and has been applied some to continuous control of ships and submarines. It still holds promise for a variety of future applications. It clearly has a role in supervisory control as a class of display aids.

2. Predictor Display for Remote Vehicle Control

When there is significant transmission delay (say more than 0.5 seconds) and slow frame-rate (say less than one frame per four seconds) a predictor display can be useful. Both of the latter conditions are likely to be present with long distance acoustic communication.

A random terrain was generated and displayed in perspective, updated every 8 seconds (Figure 33). A predictor symbol appeared on the terrain, continuously changing as the experimental subject controlled the motion of the vehicle, through a one-second time delay. Front-back velocity control was accomplished through corresponding position adjustment of a joystick, and turn rate by the left-right position of the joystick. Also superposed on the static terrain picture was a prediction of the viewpoint for the next static picture, and an outline of its field of view. This reduced the otherwise considerable confusion about how the static picture changed from one frame to the next, and served as a guide for keeping the vehicle within the available field of view. By use of the above two display symbols together, relative to the periodically updated static (but always out of date) terrain picture, subjects could maintain speed with essentially continuous control. By contrast, without the predictor they could only move extremely slowly without going unstable.

PREDICTOR DISPLAY

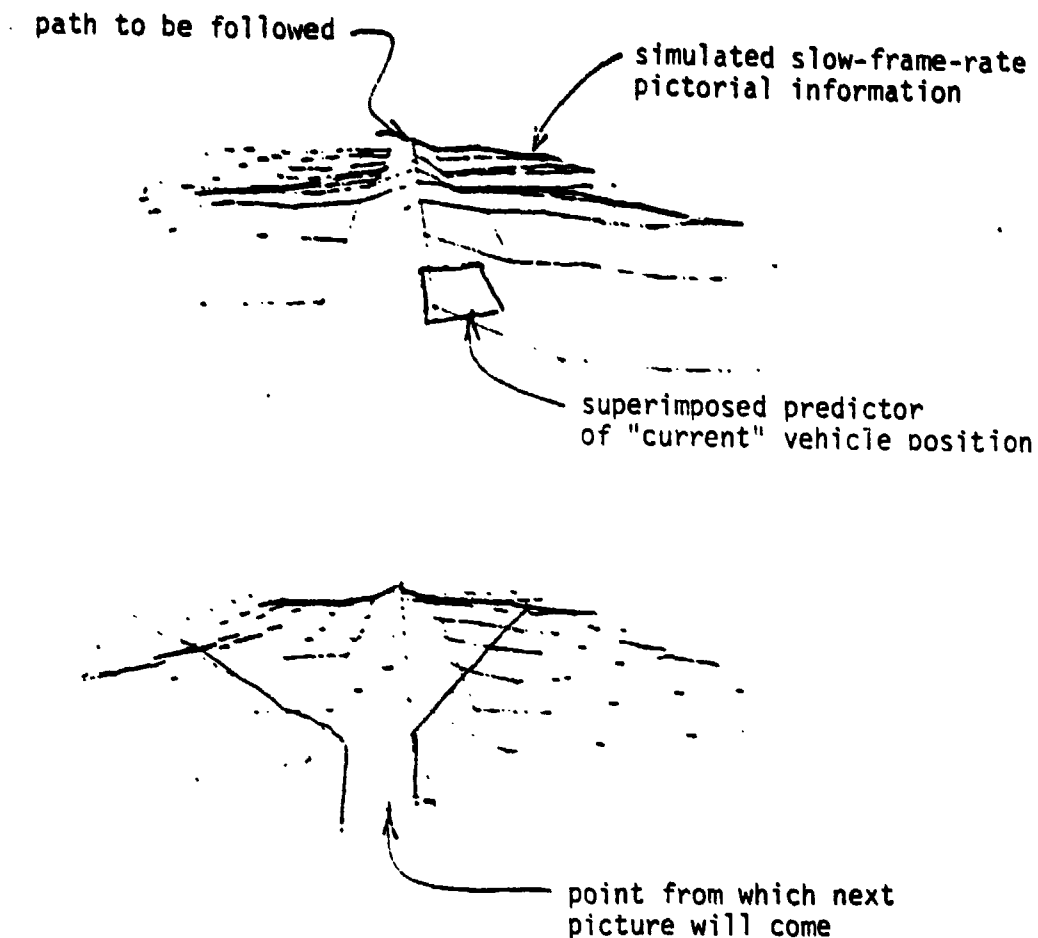


Figure 33 Simulation experiments with predictor displays (from Sheridan and Verclank, 1978). Slow-frame-rate pictures (8 seconds per frame) were simulated by computer-displayed terrain. The path to be followed was a ridge in the terrain. A moving predictor symbol (perspective square) was superimposed on the static picture of terrain. The point from which the next picture was taken was indicated with a "table" (square with four legs) and the field of view was shown with dotted lines.

E. Computer-Graphic Simulation Aids For Failure Detection/Location In Dynamic Processes

* This section is based on Tsach, Sheridan and Tzelgov [1982]

1. Introduction

The Failure Detection and Location System (FDLS) is designed to aid the operator to detect and locate failures in real-time in systems such as: power plants (either fossil or nuclear), chemical plants, airplanes, ships, etc. In these systems power is transferred from available sources to locations where it is needed and converted into a desired effective rate of work. In the normal operational mode, the power transferred from subsystem i to subsystem j is P_{ij}^n . A deviation of the measured system power, P_{ij} , from its normal value P_{ij}^n (even though both may be changing with time) indicates that the system has failed. Hence, a system failure is defined to be a process which causes the transferred power, P_{ij} , to deviate from its normal value, P_{ij}^n . Examining P_{ij} and comparing it to P_{ij}^n is the basis of the FDLS method.

It is desired to be able to detect system failures during steady state as well as transient operational modes. As a result, P_{ij}^n has to reflect the system dynamics. This can be done by constructing a dynamical model of the normal operation mode of the system. This model, implemented on a computer in real-time, will generate the desired reference power, P_{ij}^n , which is compared to the measured system power P_{ij} , as shown in Figure 34.

A simple comparison of real system to model turns out to be limited in its capabilities to detect and locate system failures. If the model is reliable, an unacceptable disparity between P_{ij} and P_{ij}^n indicates the system has failed. Yet, this comparison does not indicate whether the cause is in the i -th, j -th or neither subsystem. Namely, a system failure is detectable by this means, but its cause cannot be located. Furthermore, in the case of multiple system failures, P_{ij} might not deviate from its normal value, P_{ij}^n , although the system has failed. This happens, for example, when multiple failure results in increased effort and decreased flow such that their product is unchanged. As a result, the power comparison test will not detect these kinds of multiple failures.

2. The Model-Based FDLS

To avoid the limitations of the system-model setup shown in Figure 34, one can modify this setup in a very special way. The power transferred from subsystem 1 to subsystem 2, P_{12} , is measured and its effort and flow variables are e and f . A model of each of the two subsystems (submodel 1 and submodel 2) is constructed. According to the causalities of these submodels, the measured system effort and flow variables are used as the model inputs.

In Figure 35 it is assumed that the causalities of the submodels are such that e_s and f_s are the input variables of submodels 1 and 2, respectively. The complementary variables, f_m and e_m , calculated by the model are the model output variables. Hence, the input to each submodel is a measured system variable, the covariable of which is calculated by the

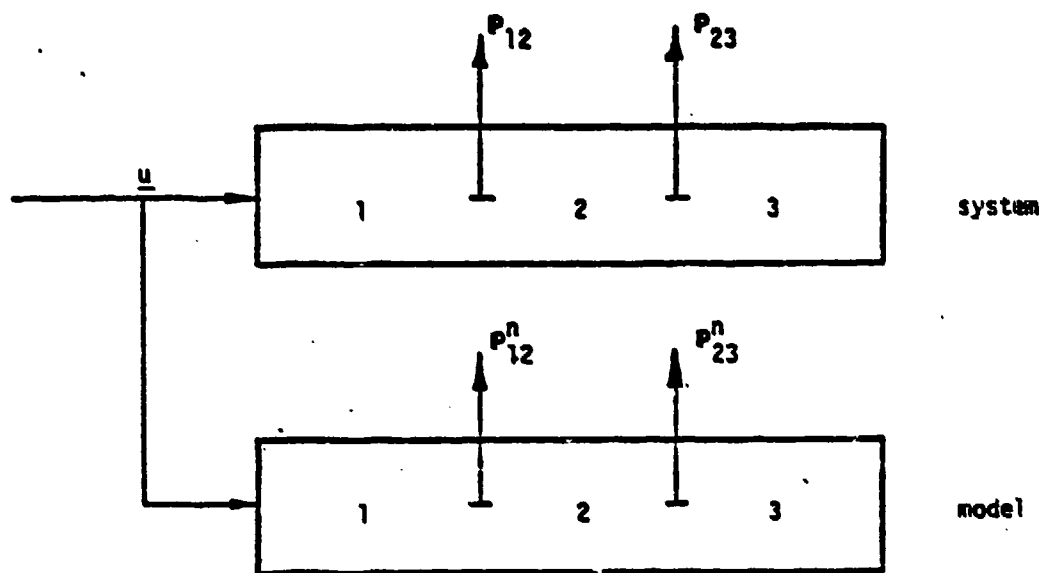


Figure 34 Comparing the system transferred power values, P_{12} and P_{23} , with their model values, P_{12}^n and P_{23}^n . \underline{u} is a vector of the system's input variables.

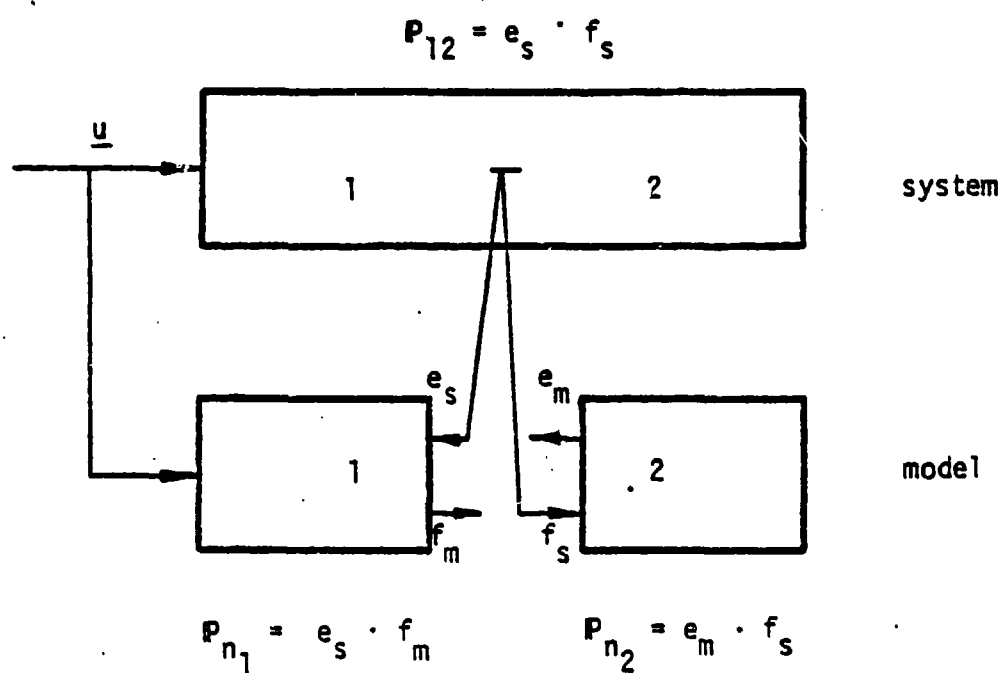


Figure 35 FDLM system-model setup where the system power, P_{12} , is compared with its model values, P_{n1} and P_{n2} .

submodel. The product of the input and the output of each submodel is the submodel calculated power. The power values of submodel 1 and 2 are:

$$P_{n1} = e_s f_m$$

$$P_{n2} = e_m f_s$$

By comparing these values with the system power P_{12} , system failures can be detected and the locations of their causes can be identified. If the model is reliable and P_{n1} differs from P_{12} , the cause of the failure is located in subsystem 1, and if P_{n2} differs from P_{12} , the cause is located in subsystem 2. A multiple system failure causes both P_{n1} and P_{n2} to differ from P_{12} . Thus the modified system-model setup shown in Figure 35 is capable of both detecting multiple system failures and locating their causes.

A comparison of the system and model power values is equivalent to the comparison of the corresponding effort or flow variables. The comparison of f_m and f_s (Figure 35) is equivalent to the comparison of P_{n1} and P_{12} , and e_m and e_s is equivalent to P_{n2} and P_{12} . In most applications the effort and flow comparison is advantageous. However, in electrical alternating current (AC) applications it is better to compare power. This is true since in AC applications both e and f variables are continuously changing rapidly, whereas power is relatively constant.

In order to refine the location of the failure causes, one has to divide the system into smaller subsystems, and to perform a number of comparison tests as described in Figure 35. These tests can be performed either simultaneously (as shown in Figure 36) or sequentially (as shown in Figure 37). The trade-off between the simultaneous and sequential tests is discussed in Tsach [1982]. There a discussion appears on how to deal with the problem of performing the detection/location quickly, before the e and f variables have wandered too far out of their normal range and the submodels are no longer valid. It is also shown how to apply the FDLs to systems which are nonlinear, not-simply-connected, or thermofluidic.

The next section deals with how the measured discrepancies between actual system and model should be displayed to the human operator, and what data processing aids can be provided.

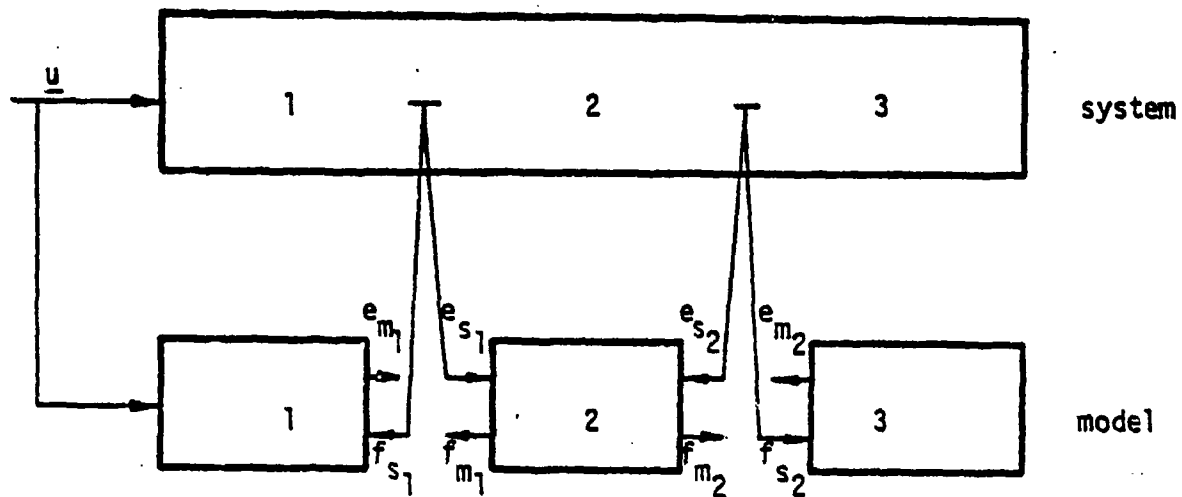


Figure 36 FDLM simultaneous mode. All the available system measured variables are utilized simultaneously and the model is disaggregated into small submodels.

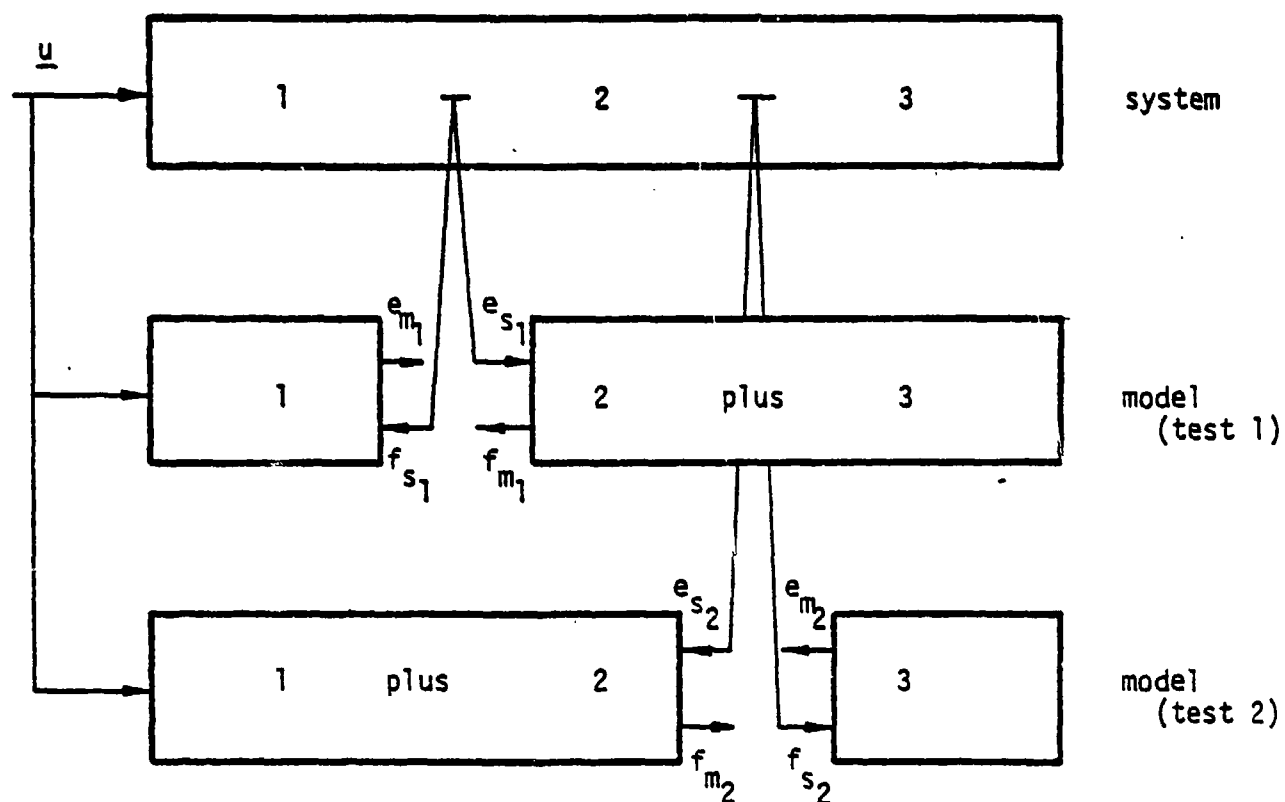


Figure 37 FDLM sequential mode. The system measured outputs of each point are utilized sequentially (test 1, test 2 ...) to pinpoint the system failure.

F. Raw Versus Processed Data For Failure Detection/Location*

* This section is predicated on Tzelgov, Tsach and Sheridan [1983] and Tsach [1982]

1. Introduction

One approach to detecting failures in dynamic systems is to run a computer model in parallel to the real system, as described in Section III-E above. In such cases the operator is expected to monitor the discrepancy and decide when it is sufficiently large to warrant concern and/or is not caused by known factors considered to be non-failures (e.g., equipment "locked out" for maintenance).

The question is then, how these variables should be presented to the operator. Should he observe the raw signals coming from real system and model, or should he observe processed information, where averaging and prejudgement is already made (suggested, or both?). This experiment was set up to answer that question.

2. Bayesian Decision Aiding

Assume that both the system and the model outputs have Gaussian distributed noise components. It follows that under normal operation mode, D (the difference between the system and the model) is a Gaussian variable with $E(D)$ and variance σ_D^2 , that can be estimated by taking measurement under nonfailure steady state conditions. An experienced operator, who has been monitoring a model system display, should be able to state the maximal model-system difference that might be expected under nonfailure conditions. has no idea whatsoever which of the two hypotheses is correct, i.e., his prior distribution is diffuse. Now a sample of size n is taken and the mean difference \bar{d} is computed. Since the sampled process is Gaussian, the posterior distribution is Gaussian as well. Parameters \bar{d} and σ_D^2/n . $p(H1)$ can be estimated as the area under the Gaussian curve prespecified by this hypothesis, and $p(H2)$ is simply $1-p(H1)$. When the next sample is taken, the previous posterior distribution may serve as prior distribution which will result in combining the information from the two samples (see [Winkler and Hayes, 1970] for details).

Note that when much data indicating no failure has been accumulated, $p(H2)$ is almost zero. If a failure appears at this stage, much additional data may be required to overcome the certainty in the nonfailure hypothesis. In order to overcome this difficulty we decided to accumulate information only if $p(H2)$ exceeds 0.5, and to assume diffuse priors otherwise. This solution is similar to resetting the decision function to zero in Gai and Curry's [1976] model of the operator.

The effect of aiding the operator in information gathering by providing him with smoothed data and the effect of aiding him in the decision making stage were evaluated independently in an experiment described below. In particular, we compared performance in the case of a raw system-model display to the performance when either the model only or both the model and the system outputs were smoothed by averaging. Such analog information

appeared either by itself or was coupled with some digital information related to failure probability. Under one condition this failure probability information was presented in terms of the failure/nonfailure odds, i.e. the ratio $\Omega = p(H2)/p(H1)$. We preferred this odds ratio over simple failure probability because of its higher sensitivity, and because evidence exists that in some tasks it results in less conservative behavior [Edwards 1963]. In another condition, whenever the probability of a failure exceeded 0.5 a "time flag," showing for how long this situation has lasted, was displayed.

An additional factor varied in the experiment was whether the compared variables (system and model) are state variables or not. Tsach [1982] has shown that when state variables (i.e., integrators) are used, the variables are automatically smooth and the resulting model of the outputs appears almost noiseless.

Under such conditions, the gathering and integrating of information can be done rather easily and, as a consequence, no effect of added smoothing should be expected when state variables (i.e., integrators) are used for comparison. However, it is still a question whether, under these almost perfect conditions for information gathering, information about failure probability will enhance performance. A positive answer to this question will help validate the two stage information processing by the human operator implied in this work.

3. Experiment

Apparatus - A "prototype" of a hydraulic-mechanical-electrical system was implemented on a PDP-11 computer (see Figure 38). The simulation has generated either the torque or the angular velocity of the generator's shaft as output variables. Gaussian noise with expectation zero and standard deviations of 0.55 of the measurement unit was added to the measured outputs.

A computerized model of the system was implemented in parallel to the prototype simulation. Failures were introduced by introducing a 10 percent disparity between the outputs of the prototype and its model. Because it is a state variable (i.e., follows an integrator in a dynamic sense) the model's torque was noiseless. However, when a failure was introduced, the model's torque diverged from its prototype only after 12-15 seconds, clearly a less desirable feature. On the other hand, both the prototype and its model angular velocity were noisy, but in case of failure, their values diverged almost immediately.

Experimental Design - The analog information appeared in a running window centered on the Megatek 7000 Display monitor. The digital information appeared above the running window (see Figure 39).

Smoothing was achieved by connecting the means of samples of size 15 taken in equal time intervals each 1.5 seconds. The same samples were used for estimation of failure probabilities. The nonfailure hypothesis was defined independently for the "state" and non-state" conditions.

The experiment was run in two sessions. Each experimental session was further divided into four blocks, one for each type of digital

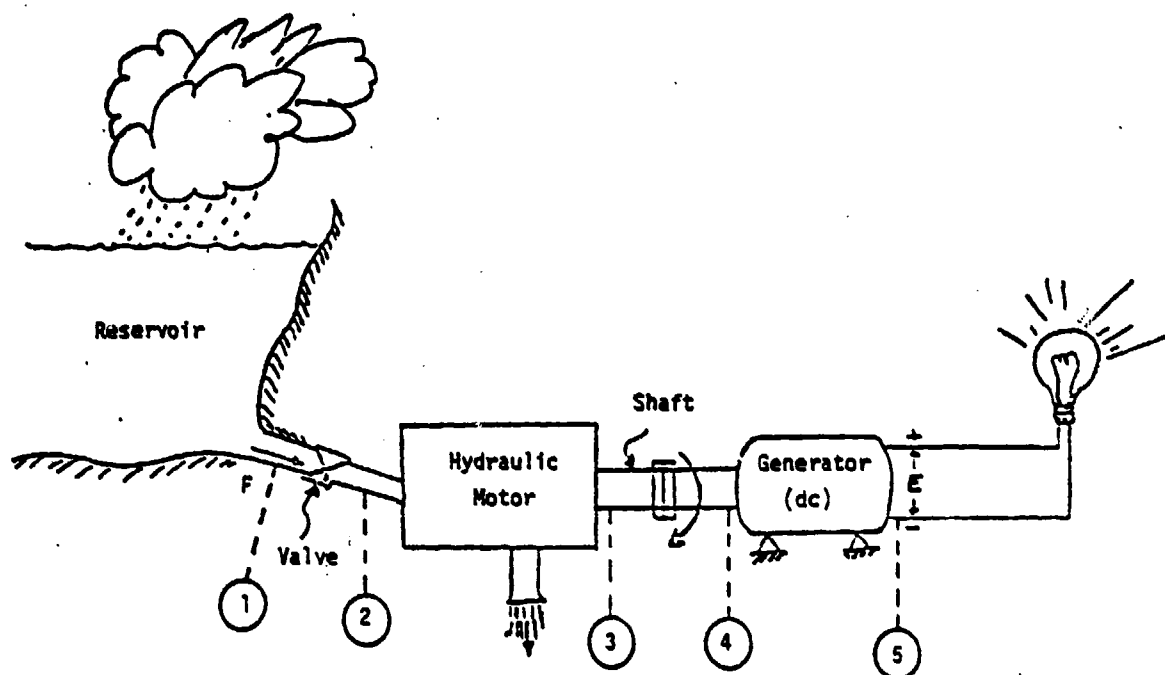
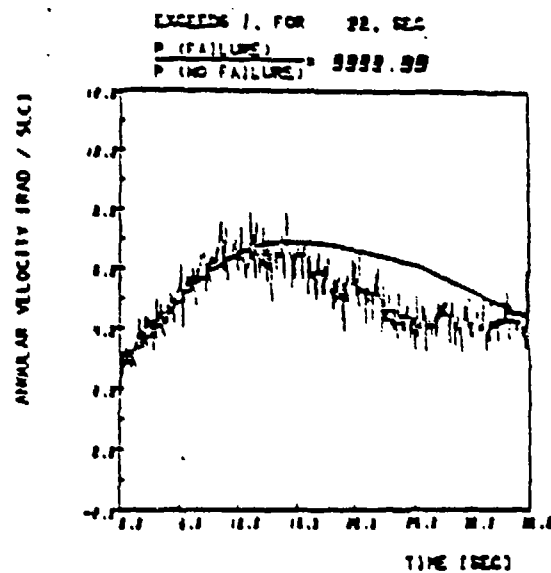


Figure 38 A description of an hydraulic-mechanical-electrical system. There are 5 measurement points available (Points 1-5).

PRACTICE
PRESS START BUTTON TO CONTINUE



RAW SYSTEM & MODEL OUTPUTS

DIS:1A. PROBABILITY RATIO AND THE TIME DURING WHICH IT EXCEEDS 1

Figure 39

information. Each block contained 60 experimental trials, 20 in each mode of digital information, 10 of them with a failure. The times at which the failure could be seen were of course dependent on the time constants of the monitored variables and were longer in the state-variable-included condition. The trials in each block were randomly ordered, by using different random orders for different blocks. The order of blocks within sessions and the order of trials within blocks were counterbalanced across subjects by a Greco-Latin square. The order of sessions was balanced across subjects.

The experiment may be described as $2 \times 2 \times 3 \times 4$ factorial design; the factor being Order of Session (state variable first vs. non-state variable first), Type of Session (state variable included vs. not), Analog Information (raw, means for model only, means for both model and the system) and Digital Information (none, odds ratio, time flag, time flag and odds ratio).

Procedure - The subjects participating in the experiment were asked to monitor in real time the system and model outputs, and by using all the information available on the CRT display to decide if the system in the monitored trial had failed. The different kinds of information that might be available for decision and the definition of failure in terms of the minimum system-model discrepancy were described and explained to the subjects in the beginning of each session. The subjects were also told that the failure should be seen on the CRT almost immediately (no-state-variable condition) or only after 15 seconds (state-variable condition). They were asked to indicate failures by pressing an "alarm" button. No overt response was required for non-failure decisions. Accuracy and speed of response were said to be of equal importance.

A session began with 20 practice trials in which only analog information was displayed. Furthermore, each block was preceded by 10 practice trials. The practice trials were followed by accuracy feedback. No feedback was given in the experimental trials. Each session lasted for about 2.5 hours.

Before each trial, an empty window and a description of the information to be available in the trial to follow were displayed on the computer. Thus, the subject knew in advance what kind of analog (as well as digital) information to expect. Subjects started a trial by pressing a start button. The trial terminated when the subject "detected a failure," i.e., pressed the alarm button. When no failure was detected a trial lasted for 35 seconds.

Subjects - Eight paid volunteers, all of them students of engineering at MIT with some background in control theory, participated in the experiment.

Results - Accuracy data appear in Table 4. As can be seen there are almost no errors in the state-variable-included case. However, in the no-state-variable-included condition, when both the system and the model are noisy and where the only available information is raw output, the proportion of error is higher, both in terms of misses and false alarms. As can be seen from Table 4 smoothing the outputs or providing the operator with failure probability estimations results in rather dramatic improvement of accuracy.

Table 4

Percentages of False Alarms (FA) and Misses (MS)

	Digital information	No Digital inf.		Ω		Time flag		time Ω & flag	
		FA	MS	FA	MS	FA	MS	FA	MS
Smoothing									
	Raw system & model outputs	13.75	11.25	0.0	0.0	3.75	0.0	0.0	0.0
No-state- variable included	Raw system & smoothed model outputs	5.0	1.25	0.0	0.0	1.2	0.0	2.5	1.25
	Smoothed system & model outputs	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	Raw system & model outputs	0.0	0.0	1.25	0.0	0.0	0.0	0.0	0.0
State- variable included	Raw system & smoothed model outputs	1.25	1.25	1.25	0.0	1.25	0.0	1.25	0.0
	Smoothed system & model outputs	2.5	0.0	0.0	0.0	2.5	0.0	2.5	0.0

Mean detection times for correctly detecting failures were subjected to a four way analysis of variance with order of session as a between-subjects factor and type of session, analog information, and digital information as within-subjects factors. It was found that detection times under the state-variable-included condition were much longer than reaction times in the no-state-variable-included condition [$F(1,6) = 75.99, p < 0.001$]. Differences in digital as well as in analog information were both found to affect performance [$F(3,18) = 5.02, p < 0.02, F(2,12) = 9.97, p < 0.01$ respectively]. However, analog information interacted with type of session [$F(2,12) = 13.74, p < 0.001$]. Moreover, a second order interaction among analog information, type of session and order of sessions was also significant [$F(2,12) = 5.26, p < 0.025$].

Newman-Keuls tests revealed that time flags do not result in hastening decision time over the no digital information conditions. However, the detection under each of these two conditions was found to be significantly slower than either in the odds ratio condition or in the odds ratio and time flag condition. The detection times in these two conditions did not differ significantly from each other. The above results are summarized in Figure 40.

The second order interaction between analog information, type of session, and order of sessions was further analyzed by testing the differences between the different types of analog information with the Newman-Keuls method for each combination of the two other interacting factors separately. As can be seen from Table 5 and as expected, smoothing had no effect in the state-variable-included condition, whereas otherwise smoothing of both the system and the model results in somewhat better performance.

4. Discussion And Conclusions

When the operator is required to detect system-model differences he is faced with problems of both information gathering and decision making. The main difficulty in information gathering is the distinction between the two outputs. It was found that such a distinction could be made easily when the model was basically smooth, i.e., when the comparison was made at state-variable points, just after integrators in the dynamics.

When both outputs were noisy, we could improve performance by smoothing the outputs of both the system and the model. However, there are good reasons for smoothing the model only. First of all, the operator should have the possibility to learn the noise characteristics of the system. Moreover, it may be important for the operator to know which of the two displayed outputs is the system and which is the model. Accordingly, we propose to leave the system unsmoothed. The efficiency of a smoothed-model raw-system display seems even more promising in the light of the expected improvement in dealing with noisy outputs as a function of practice, suggested by our data.

It was evident that the operators were aided at the decision making stage by providing them with the failure odds, which also are hypothesized to be helpful in detecting variance failures. However, in order to use such a device, the nonfailure hypothesis should be very carefully defined, and operators should be trained how to define failures in systems by providing them with feedback about the implications of their definitions.

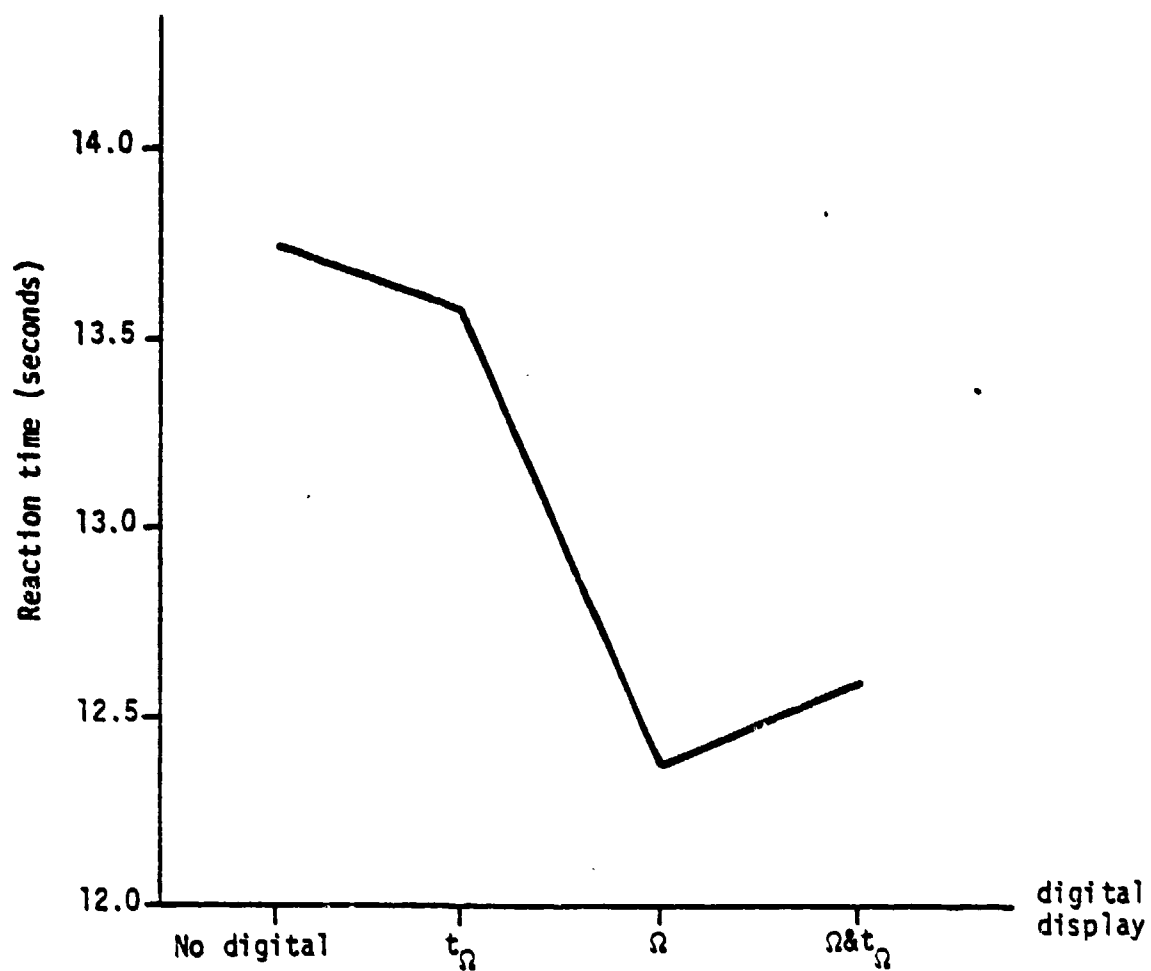


Figure 40 The operators' reaction times versus the different digital information displays.

Table 5

Mean Detection Time (in seconds) as a Function of
Session Type, Session Order, and Analog Information

First session	State-variable included			No state-variable included		
	Raw outputs	Model smoothed	Model & system smoothed	Raw outputs	Model smoothed	Model & system smoothed
State- variable included	16.76	17.15	16.83	8.00	8.97	7.54
No state- variable included	16.32	16.97	16.21	11.54	11.23	9.16

We believe that simple digital information such as failure odds may be used efficiently not only for detecting failures but also to pinpoint the faulty component. In such cases, the operator may be required to monitor quite a large number of system-model comparisons. A good way to do it, without crowding the display too much, is to provide the operator with the failure odds ratio for each comparison. When the operator is alerted by a sudden increase in one of the displayed odds ratios, he can check his suspicions by getting the relevant analog information.

A failure detection and location system (FDLS) using a man machine interface designed along these lines is being developed in our laboratory and seems quite promising.

CONCLUSIONS

A number of experiments have been described to illustrate how supervisory control can be implemented in the context of controlling remote manipulators and dynamic systems, including detecting and locating failures. The experiments cited were clustered in two groups, the first relating primarily to command -- "effectors" (human operator to computer to remote system), the second relating primarily to sensing or "affectors" (remote system to computer to human operator). In both cases the computer was providing essential mediation, much as a staff provides mediation between the top "boss" and the bottom level of workers in a large organization.

It becomes increasingly clear that supervisory control is not one function but many, and we may use further the analogy of functions served by middle-men in human organizations. When serving in the "effector" category the high level functionaries may concern themselves with translating the boss's orders into detailed instructions which are consistent with general policy (initial conditions) and recent orders, plus some forecasts (internal models) run by support staff. Lower level effector functionaries may receive these more specific orders and control the workers to put them into effect, making use of direct feedback from the workers. Such computer functions were embodied in the supervisory command systems described in Sections 2.1 and 2.2. Some other staffers may specialize in helping the boss formulate company policy by pointing out environmental facts (analogous to the "pointing" techniques in Section 2.3). Some middle managers are assigned to nulling disturbances that arise (as in Section 2.4) while some management advisors are concerned with how various personnel should allocate their time (as related to Section 2.5).

In the "affecter" category many of these staff mediators may be concerned about how and what to sample in the environment, much as economists, accountants and market researchers would (as with Section 3.1 and 3.3), whereas other mediators may use the sampled data plus a priori parametric information to run simulations of the situation to test and plan future actions (Sections 3.2 and 3.4). In some cases, the mediators may exercise such models as aids to detecting and diagnosing failures within company operations (Sections 3.5 and 3.6).

The main point is that in supervisory control the computer does not provide one function or one form of mediation, but many - at different places in the system, at different times or under different circumstances. This we believe will become evident with the various application of supervisory control - e.g., aircraft piloting and traffic control and office automation - as well as remote manipulation and process control. Nevertheless, though the computer takes on a diversity of functions, from the human operator's viewpoint the change to being a supervisor is always a change from continuous and direct sensing and control to indirect or somewhat remote control. The change means observing more integrated displays and issuing subgoal or conditional commands, all at a higher level than with continuous direct control.

The motivations for going to supervisory control are also many, as was implied in earlier discussion. They may be technical constraints on

bandwidth, transmission time delay, or simple inability to provide sufficiently fast or accurate control signals in the direct manual control mode. Or it may be that the human operator is too busy or too fatigued or too bored to remain consistently in the control loop. In any such case the test of supervisory control is whether it works better.

It is proving difficult to determine what is "optimal" supervisory control. For one thing in supervisory control one of the operator's primary tasks is to set and modify subgoals and criteria; the "objective function" is not fixed. A second factor is the difficulty of modeling and experimenting on supervisory systems because of the inherent cost and complexity: it is not feasible to vary all parameters independently to find the best mix; there are simply too many parameters and the cost of changing is too great. Just the collection of data in the supervisory case proves difficult since the measures of human performance now have to do with how the operator communicated with the computer, what various displays he observed, what concatenations of commands he issued - each event likely to be different from the last. As supervisory control systems become more sophisticated it becomes less and less likely that any one stimulus or response situation will be repeated.

So we seem to have to abandon our simple behavioral analytical models of the operator, where we have statistical confidence in few parameters by dint of many repetitions in well controlled situations, with few variables changing at a time. We feel a push to adapt a more holistic/synthetic approach to engineering man-machine systems.

In the pursuit of experimental control and clean straightforward science we cannot simply retreat to simpler manual control. Supervisory control is here, it works, it will be used and demanded to be made better.

REFERENCES

Brooks, T.L., 1979, "SUPERMAN: A System for Supervisory Manipulation and the Study of Human/Computer Interactions", SM Thesis, MIT Man-Machine Systems Laboratory Report.

Brooks, T.L. and Sheridan T.B., 1980, Annual Manual, Experimental Evaluation of the Concept of Supervisory Manipulation, Proc., 16th Annual Conference on Manual Control, Cambridge, MA.

Deghues, B.J., 1980, Operator - Adjustable Frame-Rate, Resolution and Gray Scale Tradeoff in Fixed-Bandwidth Remote Manipulation Control, SM Thesis, MIT Man-Machine Systems Laboratory Report.

Edwards, W., 1968, Conservatism in Human Information Processing. In B. Kleinmütz (Ed.) Formal Representation in Human Judgement, N.Y., Wiley.

Fyler, D., 1981, "Computer Graphic Representation of Remote Environment Using Position Tactile Sensors", SM Thesis, MIT Man-Machine Systems Laboratory Report.

Gai, E. and Curry R., 1976, A Model of the Human Observer in Failure Detection Tasks, IEEE Trans. System Man and Cybernetics, 6, 85-94.

Hirabayashi, H., 1981, Supervisory Manipulation for Assembling Mechanical Parts While Compensating for Relative Motion, SM Thesis, MIT Man-Machine Systems Laboratory Report.

McRuer, D.T., and Krendel, E.S., 1957, Dynamic Response of Human Operators, WADC-TR-56-524, U.S. Air Force.

Ranadive, V., 1979, Video Resolution, Frame Rate and Gray Scale Tradeoffs under Limited Bandwidth for Undersea Teleoperation, SM Thesis, MIT Man-Machine Systems Laboratory Report.

Ranadive and Sheridan, 1981, Video Frame Rate, Resolution and Grayscale Tradeoffs for Undersea Telem manipulator Control, Paper submitted to IEEE Trans Systems Man and Cybernetics.

Sheridan, T.B., 1982, Supervisory Control: Problems Theory and Experiment for Application to Human-Computer Interaction in Undersea Remote Systems, MIT Man-Machine Systems Laboratory Report.

Sheridan, T.B., and Ferrell, W.R., 1974, Man-Machine Systems, MIT Press.

Sheridan, T.B., and Verplank, W.I., 1978, Human and Computer Control of Undersea Teleoperators, MIT Man-Machine System Laboratory Report.

Tani, K., 1980, Supervisory Control of Remote Manipulation with Compensation for Moving Targets, MIT Man-Machine Systems Laboratory Report.

Tsach, U., 1982, Failure Detection and Location Method (FDLM), Ph.D.

Thesis, MIT Man-Machine Systems Laboratory Report.

Tsach, Sheridan and Tzelgov, 1982, A New Method for Failure Detection and Location in Complex Dynamic Systems, Proceedings 1982 American Control Conference, June, Arlington, VA.

Tzelgov, Tsach and Sheridan, 1983, Effects of Indicating Failure Odds and Smoothed Outputs on Human Failure Detection in Dynamic Systems, Paper submitted to Human Factors.

Tzelgov, Yoerger and Sheridan, 1983, Factors Affecting the Angular Accuracy of Pointing with a Master Slave Manipulator, MIT Man-Machine Systems Laboratory Report.

Winey, C.M., 1981, Computer Simulated Visual and Tactile Feedback as an Aid to Manipulator and Vehicle Control, SM Thesis, MIT Man-Machine Systems Laboratory Report.

Winey C.M. and Sheridan, T.B., 1983, Computer Simulated Feedback as an Aid to Manipulator and Vehicle Control, Paper submitted to IEEE Trans. on Systems, Man and Cybernetics.

Winkler R.L. and Hayes, W.L., 1970, Statistics, Probability Inference and Decision, N.Y., Winston.

Wood, W. T., 1982, The Use of Machine Aids in Dynamic Multi-Task Environments: A comparison of an Optimal Model to Human Behavior, MIT SM Thesis, MIT Man-Machine Systems Laboratory Report.

Yoerger, D., 1982, Supervisory Control of Underwater Telem manipulators: Design and Experiment, Ph.D. Thesis, MIT Man-Machine Systems Laboratory Report.

Yoerger, D. and Sheridan, T.B., 1983, Supervisory Control Improves Performance for Underwater Telem manipulators, Marine Technology Society Meeting on Remotely Operated Vehicles, March, San Diego.

DISTRIBUTION LIST

Capt. Paul R. Chatelier
Office of the Deputy Under
Secretary of Defense
OUSDRE (E&LS)
Pentagon, Room 3D129
Washington, D.C. 20301

Engineering Psychology Programs
Code 422
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (2cys)

Undersea Technology Programs
Code 220
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Dr. Eugene Silva
Ocean Sciences, Code 421
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Communication & Computer
Technology Programs
Code 240
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Manpower, Personnel,
and Training Programs
Code 270
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Information Sciences Division
Code 433
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Special Assistant for Marine
Corps Matters
Code 100M
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Eastern Regional Office
ATTN: Dr. J. Lester
Building 114, Section D
666 Summer Street
Boston, MA 02210

Dr. Perry Alers
Marine Systems Branch
Code 5821
Naval Research Laboratory
Washington, D.C. 20375

Director
Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375

Office of the Chief of
Naval Operations
Deep Submergence Systems Div.
OP-231
Washington, D.C. 20350

Dr. Robert G. Smith
Office of the Chief of Naval
Operations, OP 987H
Personnel Logistics Plans
Washington, D.C. 20350

Dr. Andreas B. Rechnitzer
Office of the Chief of Naval
Operations, OP 952F
Naval Oceanography Division
Washington, D.C. 20350

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, FL 32813

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, Ca 93940

Human Factors Technology Administrator
Office of Naval Technology
Code MAT 0722
800 North Quincy Street
Arlington, VA 22217

Dr. Robert Fleming
Human Engineering Branch
Code 8231
Naval Ocean Systems Center
San Diego, CA 92152

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

Dr. Andreas B. Rechnitzer
Office of the Chief of
Naval Operations, OP 952F
Naval Oceanography Division
Washington, DC 20350

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, FL 32813

Mr. John Freund
NAVSEA 05R2
Naval Sea Systems Command
Washington, DC 20362

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, CA 93940

Mr. John Freund
NAVSEA 05R2
Naval Sea Systems Command
Washington, D.C. 20365

Mr. H. Talkington
Ocean Engineering Department
Naval Ocean Systems Center
San Diego, CA 92152

Mr. Paul Heckman
Naval Ocean Systems Center
San Diego, CA 92152

Dr. Ross L. Pepper
Naval Ocean Systems Center
Hawaii Laboratory
P.O. Box 997
Kailua, HI 96734

Mr. Phillip Andrews
Naval Sea Systems Command
NAVSEA 0541
Washington, D.C. 20362

Mr. Milon Essoglou
Naval Facilities Eng. Command
R&D Plans and Programs
Code 03T
Hoffman Building II
Alexandria, VA 22332

CDR Robert Biersner
Naval Medical R&D Command
Code 44
Naval Medical Center
Bethesda, MD 20014

Dr. Arthur Bachrach
Behavioral Sciences Department
Naval Medical Research
Institute
Bethesda, MD 20014

Mr. H. Talkington
Ocean Engineering Department
Naval Ocean Systems Center
San Diego, CA 92152

Human Factors Technology Administrator
Office of Naval Technology
Code MAT 0722
800 N. Quincy Street
Arlington, VA 22217

Dr. George Moeller
Human Factors Engineering Branch
Submarine Medical Research Lab
Naval Submarine Base
Groton, CT 06340

Dr. Robert Blanchard
Navy Personnel Research and
Development Center
Command and Support Systems
San Diego, CA 92152

Mr. J. Williams
Department of Environmental Sciences
U.S. Naval Academy
Annapolis, MD 21402

Mr. John Quirk
Naval Coastal Systems Laboratory
Code 712
Panama City, FL 32401

Technical Director
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Technical Director
U.S. Army Human Engineering Labs
Aberdeen Proving Ground, MD 21005

Dr. T. B. Sheridan
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139

CDR Thomas Berghage
Naval Health Research Center
San Diego, CA 92152

U.S. Air Force Office of
of Scientific Research
Life Sciences Directorate, NL
Bolling Air Force Base
Washington, DC 20332

Chief, Systems Eng. Branch
Human Engineering Division
USAF AMRL/HES
Wright-Patterson AFB, OH 45433

Defense Technical Information
Center
Cameron Station, Bldg. 5
Alexandria, VA 22314 (12 cys)

Dr. Craig Fields
Director, System Sciences
Office
Defense Advanced Research
Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Dr. M. Montemerlo
Human Factors & Simulation
Technology, RTE-6
NASA HQS
Washington, DC 20546

Dr. J. Miller
Florida Institute of
Oceanography
University of South Florida
St. Petersburg, FL 33701

Dr. Babur M. Pulat
Dept. of Industrial Engineering
North Carolina A&T State Univ.
Greensboro, NC 27411

Dr. Richard W. Pew
Information Sciences Division
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Harry Snyder
 Department of Industrial Engineering
 Virginia Polytechnic Institute
 State University
 Blacksburg, VA 24061

Dr. Robert T. Hennessey
 NAS - National Research Council
 2101 Constitution Ave., N.W.
 Washington, DC 20418

Dr. Amos Freedy
 Perceptronic
 6271 Variel Avenue
 Woodland Hills, CA 91364

Dr. James H. Howard, Jr.
 Department of Psychology
 Catholic University
 Washington, DC 20064

Dr. Christopher Wickens
 University of Illinois
 Department of Psychology
 Urbana, IL 61801

Dr. Edward R. Jones
 Chief, Human Factors Engineering
 McDonnell-Douglas
 Astronautics Company
 St. Louis Division
 Box 516
 St. Louis, MO 63166

Dr. Douglas Towne
 University of Southern CA
 Behavioral Technology Lab.
 3716 S. Hope Street
 Los Angeles, CA 90007

Dr. A. K. Bajcsy
 Jet Propulsion Lab.
 California Institute of Tech.
 Pasadena, CA 91125

Dr. William R. Uttal
 Institute for Social Research
 University of Michigan
 Ann Arbor, MI 48109

Dr. Robert D. Ballard
 Dept. of Ocean Engineering
 Woods Hole Oceanographic
 Institution
 Woods Hole, MA 02543